

Towards a better approximation for SPARSEST CUT?

Sanjeev Arora*

Rong Ge†

Ali Kemal Sinop‡

April 12, 2013

Abstract

We give a new $(1 + \epsilon)$ -approximation for SPARSEST CUT problem on graphs where small sets expand significantly more than the sparsest cut (sets of size n/r expand by a factor $\sqrt{\log n \log r}$ bigger, for some small r ; this condition holds for many natural graph families). We give two different algorithms. One involves Guruswami-Sinop rounding on the level- r Lasserre relaxation. The other is combinatorial and involves a new notion called *Small Set Expander Flows* (inspired by the *expander flows* of [ARV09]) which we show exists in the input graph. Both algorithms run in time $2^{O(r)} \text{poly}(n)$.

We also show similar approximation algorithms in graphs with genus g with an analogous local expansion condition.

This is the first algorithm we know of that achieves $(1 + \epsilon)$ -approximation on such general family of graphs.

1 Introduction

This paper concerns a new and promising analysis of Lasserre [Las02]/Parrilo [Par03] SDP relaxations for the (uniform) SPARSEST CUT problem, which gives $(1 + \epsilon)$ -approximation on several natural families of graphs. Note that Lasserre/Parillo relaxations subsume all relaxations for the problem that were previously analysed: the spectral technique of Alon-Cheeger [AM85], the LP relaxation of Leighton-Rao [LR99] with approximation ratio $O(\log n)$, and the SDP with triangle inequality of Arora, Rao, Vazirani [ARV09] with approximation ratio $O(\sqrt{\log n})$. The approximation ratio of $O(\sqrt{\log n})$ has proven resistant to improvement in almost a decade (and there is some evidence the ratio may be tight for the ARV relaxation; see Lee-Sidiropoulos [LS11]). For a few families of graphs such as graphs of constant genus, an $O(1)$ -approximation is known.

Recently, there has been increasing optimism among experts that Lasserre [Las02]/Parrilo [Par03] relaxations—which are actually a hierarchy of increasingly tighter relaxations whose r th level can be solved in $n^{O(r)}$ time—may provide better approximation algorithms for SPARSEST CUT as well as other problems such as MAX CUT and UNIQUE GAMES, and possibly even refute Khot’s unique games conjecture. For instance Barak, Raghavendra, and Steurer [BRS11], relying on the earlier

*Princeton University, Computer Science Department and Center for Computational Intractability. Email: arora@cs.princeton.edu

†Princeton University, Computer Science Department and Center for Computational Intractability. Email: rongge@cs.princeton.edu

‡Princeton University, Computer Science Department and Center for Computational Intractability. Email: asinop@cs.cmu.edu

subexponential algorithm of Arora, Barak, Steurer [ABS10], showed that Lasserre relaxations can be used to design subexponential algorithms for the UNIQUE GAMES problem. Independently, Guruswami and Sinop [GS11] gave another rounding that looks quite different but yielded very similar results. Subsequently, Barak et al. [BHK⁺12] showed that Lasserre relaxations can easily dispose of families of UNIQUE GAMES instances that seemed “difficult” for simpler SDP relaxations: many families of instances can be solved near-optimally in 4-8 rounds! This result was subsequently extended by O’Donnell and Zhou [OZ13] to “difficult” families of graphs from [DKSV06] which are integrality gaps for uniform sparsest cut and balanced separator. Of course, it is unclear whether this demonstrates the power of these relaxations, or merely the limitations of our current lowerbound approaches. Nevertheless, the rise in researchers’ hopes for better algorithms is palpable.

But the stumbling blocks in this quest are also quite clear. First, known ideas for analysing Lasserre relaxations generally require some condition on the r^{th} eigenvalue of the Laplacian for some small r , whereupon some $f(r, \epsilon)$ levels of Lasserre are shown to suffice for $(1 + \epsilon)$ -approximation. Unfortunately, many real-life graphs (eg, even the 2D-grid) do not satisfy this eigenvalue condition so new ideas seem needed.

Another stumbling block has been the inability to relate these new rounding algorithms for Lasserre relaxations to existing SDP rounding algorithms such as Goemans-Williamson and ARV. Since Lasserre relaxations greatly generalize normal SDP relaxations, one would like general purpose rounding algorithms which for small r reduce to earlier rounding algorithms. A concrete question is: does the Guruswami-Sinop rounding algorithm always give an approximation ratio as good as the ARV $\sqrt{\log n}$ for SPARSEST CUT once r is sufficiently large? This is still unclear.

The current paper makes some progress on these stumbling blocks. We show that the GS rounding algorithm achieves $(1 + \epsilon)$ -approximation for SPARSEST CUT on an interesting family graphs that are *not* small set expanders and may not have large r th eigenvalue. If ϕ_{local} denotes the minimum sparsity of sets of size n/r , and ϕ_{sparsest} the minimum sparsity among *all* sets, then we require $\phi_{\text{local}}/\phi_{\text{sparsest}} \gg \sqrt{\log n \log r}$. Note that ϕ_{local} is often larger than ϕ_{sparsest} in natural families of graphs. For example, in normalized d -dimensional $n^{1/d} \times \dots \times n^{1/d}$ -grid graphs, $\phi_{\text{sparsest}} \leq \frac{1}{dn^{1/d}}$, $\phi_{\text{local}} \gg \frac{1}{d} \left(\frac{r}{n}\right)^{1/d}$ whereas $\lambda_r \ll \frac{1}{d} \left(\frac{r}{n}\right)^{2/d}$. Note that when the condition is not met, a simple modification of our algorithm returns a subset of size n/r that has sparsity $\sqrt{\log n \log r}$ times ϕ_{sparsest} . Thus setting $r = O(1)$ one recovers the ARV bound —though the analysis of this case also uses ARV¹.

Comparison with existing work. As mentioned, earlier analyses of Lasserre relaxations require a lowerbound on the r^{th} eigenvalue of the graph: the tightest such result from [GS13] requires $\lambda_r > \phi_{\text{sparsest}}$. Efforts to get around such limitations have focused on understanding structure of graphs which *do not* satisfy the eigenvalue condition: an example is the so-called *high order Cheeger inequality* of [ABS10] (improved by Louis et al [LRTV12] and Lee et al. [LGT12]) according to which —roughly speaking—a graph with many eigenvalues close to $o(1)$ have a small nonexpanding set. In other words, the graphs are not *Small-Set Expanders*². However, there is an inherent

¹We also know how to achieve qualitatively similar results as our main result using BRS rounding + ARV ideas applied to Lasserre solutions at the expense of stricter requirements on small set expansion. However, that method seems unable to give better than $O(1)$ -approximation, whereas GS rounding is able to give $(1 + \epsilon)$.

²In fact, the unexpected appearance of Small Set Expansion (SSE) in this setting is believed to not be a fluke. It appears in the SSE conjecture of Raghavendra and Steurer [RS10] (known to imply the UGC), their “Unique games with SSE” conjecture, as well as in the known subexponential algorithms for UNIQUE GAME. Furthermore, attempts to construct difficult examples for known SDP-based algorithms also end up using graphs (such as the noisy

Cheeger-like gap (ϕ vs $\sqrt{\phi}$) between eigenvalues and expansion that seems to limit the possible improvements. Our algorithms work even without a bound on the r^{th} eigenvalue; they only need bounds on expansion (The d -dimensional grids are good examples.) Furthermore, they yield $(1+\epsilon)$ -approximation, which in context of SPARSEST CUT seems quite surprising.

Subsequent to our work and inspired by it, Gharan and Trevisan [GT13] have shown how to obtain factor $O(\sqrt{\log k})$ approximation from the basic ARV relaxation for the sparsest cut problem under local expansion or spectral conditions.

Better algorithms for bounded genus graphs. Recall that for genus g graphs there are known $O(\log g)$ -approximation algorithms for SPARSEST CUT [LS10]. We can show that GS'13 rounding gives a $(1+\epsilon)$ -approximation if $\phi_{local} \geq \Omega(\frac{\log g}{\epsilon^2})\phi_{sparsest}$. Thus for the $2D$ -grid, it implies that $O(1/\epsilon^4)$ rounds of Lasserre yield a $(1+\epsilon)$ -approximation. Again, when the local expansion condition is not satisfied our algorithm finds a witnessing small set, allowing us to recover the existing $O(\log g)$ approximation for the general case.

Combinatorial algorithm. In addition to the above Lasserre-based algorithm, we also give a new combinatorial algorithm with similar (but somewhat weaker) guarantees. This algorithm is inspired by the primal-dual algorithms for SPARSEST CUT stemming from the *expander flows* notion of ARV (see [AHK04, AK07, She09]). We introduce a new notion called *small set expander flows*: a multicommodity flow whose demand graph is an expander on small sets. Let a (r, d, β) -flow be an undirected multicommodity flow in which d units of flow is incident to each node, and the demand graph has expansion β on sets of size at most n/r (in other words, the amount of flow leaving the set S is $d\beta|S|$). We show that in every graph there is an SSE flow with $d = \Omega(\phi_{local}\sqrt{\log r}/\sqrt{\log n})$, $\beta = \Omega((\log r)^{-2})$, and this flow—or something close to it—can be found in polynomial time. Using such flows one can—with some more work—compute a $(1+\epsilon)$ -approximation to SPARSEST CUT as above.

Note that the expander flow idea of ARV was motivated by the observation that expander flows consist of a family of dual solutions to the SDP. We suspect that something analogous holds for SSE flows and the Lasserre relaxation but are unable to prove this formally. However, we can informally show a connection as follows: if a graph has a (r, d, β) -flow where

$$d\beta^2/\log r \gg \text{value of } O(r)\text{-rounds of Lasserre relaxation}$$

then the integrality gap of the Lasserre relaxation is at most $(1+o(1))$. Thus the existence of expander flows is another reason—besides the more direct rounding approach mentioned earlier—why Lasserre relaxations are near-optimal when $\phi_{local}/\phi_{global} \gg \sqrt{\log n \log r}$.

Applications to semirandom models Recently there has been interest in solving sparsest cut on semirandom models of graphs [MMV12]. In these graphs one starts with a planted sparse cut in a random graph or expander, and then an adversary is allowed to change some edges. Our work provides a new algorithm for one such model, planted combinatorial expander on regular graphs. However our results for this model are not directly comparable to the ones in [MMV12]. Our result is presented in Appendix C.

hypercube) which are small set expanders.

2 Preliminaries and Background

2.1 Expansion and Graph Laplacian

Let $G = (V, E)$ be an undirected graph with edge capacities $c_e \geq 0$ for all $e \in E$. For simplicity we assume that the input graph is regular with (normalized) degree 1, that is, for all vertices $i \in V$ $\sum_j c_{(i,j)} = 1$ (our results in Sections 3 and 4 can also be applied to irregular graphs). We always use n to denote the number of vertices in G .

The *expansion* of a set is defined as $\Phi(S) = \frac{E(S, V \setminus S)}{\min\{|S|, n - |S|\}}$, where $E(A, B) = \sum_{i \in A, j \in B} c_{(i,j)}$. The sparsity of a set $\phi(S)$ is defined as $\frac{n \cdot E(S, V \setminus S)}{|S| \cdot (n - |S|)}$. There are several problems related to sparsity of cuts:

- The *sparsest cut* of the graph is a set S that minimizes the sparsity $\Phi(S) = \frac{E(S, V \setminus S)}{|S| \cdot (n - |S|)}$. We use Φ_{sparsest} to denote its expansion and ϕ_{sparsest} to denote its sparsity.
- The *edge expansion* of a graph is a set S that minimizes the expansion $\alpha(S)$. We use Φ_{global} to denote its expansion. Notice that since we are working with regular graphs, this is also equivalent to the *graph conductance* problem.
- The *c-balanced separator* of a graph is a set S that minimizes the expansion $\Phi(S)$ among all sets of size at least cn . We use $\Phi_{\text{c-balanced}}$ to denote its expansion.

While all these problems are closely related (for example, sparsest cut and edge expansion are equivalent up to a factor of 2), we carefully differentiate between them in this paper because we are looking for $1 + \epsilon$ approximation algorithms.

We are also interested in the expansion of small sets: let $\Phi_r(G)$ be the smallest expansion of a set of size at most n/r and $\phi_r(G)$ be the smallest sparsity of a set of size at most n/r . Sometimes when r is fixed (or understood) we drop r and use Φ_{local} and ϕ_{local} instead³.

Notice that the requirement of our algorithms will have the form $\phi_{\text{local}}/\phi_{\text{global}} \gg f(n, r)$ ⁴. Since sparsity ϕ and expansion Φ are within a factor of 2 ($\Phi(S) \leq \phi(S) \leq 2\Phi(S)$), in such requirements the ratios $\phi_{\text{local}}/\phi_{\text{global}}$ and $\Phi_{\text{local}}/\Phi_{\text{global}}$ can be interchanged.

The adjacency matrix A of the graph G is a matrix whose (i, j) -th entry is equal to $c_{(i,j)}$. If $d_i = \sum_{(i,j) \in E} c_{(i,j)}$ denotes the degree of i -th vertex with D being the diagonal matrix of degrees, then the Laplacian of the graph G is defined as $L = D - A$ (for regular graph this is just $I - A$). The normalized Laplacian of the graph is defined as $\mathcal{L} = D^{-1/2} L D^{-1/2}$.

Graph Laplacians are closely related to the expansion of sets. In particular, the Rayleigh Quotient of a vector x , $R(x) = \frac{x^T L x}{x^T x}$ is exactly equal to the sparsity of a set S when x is the indicator vector of S (and S has size at most $1/2$).

We will denote by ϕ_{SDP} the optimum value of the Lasserre relaxation for SPARSEST CUT. The number of levels in the Lasserre hierarchy will be implicit in the context.

2.2 Lasserre Relaxation and GS Rounding

We will show sufficient conditions under which r rounds of Lasserre Hierarchy relaxation can be rounded to $(1 + \epsilon)$ -approximation for sparsest cut and related problems. In particular, we will

³ Φ_{local} and ϕ_{local} usually denote the optimal expansion and sparsity of sets of size at most $O(n/r)$

⁴ $f \gg g$ means $f \geq Cg$ for some large universal constant C

show that the particular rounding algorithm from [GS13] outputs such an approximation. (See Appendix A for details on the Lasserre relaxation and the GS rounding algorithm.)

In general working with Lasserre relaxations involves tedious notation involving subsets of variables and assignments to them. Luckily all that has been handled in [GS13], leaving us to work with the relatively clean (standard) SDP notation.

For the sake of simplicity, we will focus on the uniform sparsest cut problem on regular graphs. Other variants, such as edge expansion, can easily be handled by changing the objective function. Let $[x_u]_{u \in V}$ be the vectors corresponding to each node in G obtained as a solution for r -rounds of Lasserre Hierarchy relaxation. In particular, x_u 's minimize the following ratio:

$$\phi_{SDP} \triangleq \frac{\sum_{u < v} C_{uv} \|x_u - x_v\|^2}{\frac{1}{n} \sum_{u < v} \|x_u - x_v\|^2} \leq \phi_{sparsest}.$$

The denominator, whose value we will denote by ν , can also be written as:

$$\underbrace{\frac{1}{n} \sum_{u < v} \|x_u - x_v\|^2}_{\triangleq \nu} = \sum_u \left\| x_u - \frac{1}{n} \sum_v x_v \right\|^2.$$

We will shift each vector x_u by the mean:

$$X_u \triangleq x_u - \frac{1}{n} \sum_v x_v, \tag{1}$$

so that $\sum_u X_u = 0$. Note:

$$\|X_u\|^2 \leq 1.$$

We use $X = [X_u]$ to denote the matrix whose columns correspond to the vectors X_u . Since $X_u - X_v = x_u - x_v$, $X \in \ell_2^2$ (i.e. columns of matrix X satisfy the triangle inequality) and:

$$\sum_{uv} C_{uv} \|X_u - X_v\|^2 = \sum_{uv} C_{uv} \|x_u - x_v\|^2 \leq \phi_{SDP} \frac{1}{n} \sum_{u < v} \|x_u - x_v\|^2 = \phi_{SDP} \|X\|_F^2$$

where last identity follows from the fact that $\sum_u X_u = 0$. Using X , we can re-state Theorem 3.1 from [GS13] in the following way:

Theorem 2.1 (Theorem 3.1 from [GS13]). *If there exists a subset $S \in \binom{V}{r}$ with*

$$\|X_S^\perp X\|_F^2 = \sum_u \|X_S^\perp X_u\|^2 \leq \gamma \|X\|_F^2, \tag{2}$$

then the rounding algorithm from [GS13] outputs a set T such that:

$$\phi_G(T) \leq \frac{\phi_{SDP}}{1 - \gamma}.$$

Here X_S is the projection matrix onto the span of the submatrix indexed by S and X_S^\perp is the projection matrix onto the orthogonal complement of X_S 's column span.

Furthermore, the SDP solver and rounding procedure can be implemented in time $2^{O(r)} \text{poly}(n)$ using [GS12a].

3 Proof via orthogonal separators

Theorem 2.1 implies that for $(1+\epsilon)$ -approximation it suffices to show the *existence* of a small subset S of vertices such that the relative distance of all other vertices to the span of X_S is smaller than any small constant.

Theorem 3.1 (Main). *For every $\epsilon > 0$ there is a constant $C = C(\epsilon)$ such that the following is true. When all subsets of at most $2n/r$ vertices have sparsity $\phi_{\text{local}} \geq C\phi_{SDP}\sqrt{\log n \log r}$ in the graph, there exists a set S of r vertices such that $\|X_S^\perp X\|_F^2 \leq \epsilon \|X\|_F^2$. (Here $\phi_{SDP} \leq \phi_{\text{sparsest}}$ is the value of the Lasserre relaxation for $r+3$ rounds and X 's are the corresponding vectors from eq. (1).)*

This existence result will be proven using the *orthogonal separators* [CMM06] but with the modifications of Bansal et al. [BFK⁺11], which, not surprisingly, were also developed in context of algorithms for small set expansion. (We know how to give a more direct proof without using orthogonal separators but it brings in an additional factor of $\log r$ in the local expansion condition.)

Definition 3.2 (Orthogonal Separator). Let X be an ℓ_2^2 space. A distribution over subsets of X is called an m -orthogonal separator with *distortion* D , *probability scale* $\alpha > 0$ and *separation threshold* $\beta < 1$ if the following conditions hold for $S \subset X$ chosen according to this distribution.

1. For all $X_u \in X$, $\Pr[X_u \in S] = \alpha \|X_u\|^2$.
2. For all $X_u, X_v \in X$ with $\|X_u - X_v\|^2 \geq \beta \min\{\|X_u\|^2, \|X_v\|^2\}$,

$$\Pr[X_u \in S \text{ and } X_v \in S] \leq \frac{\min\{\Pr[X_u \in S], \Pr[X_v \in S]\}}{m}.$$

3. For all $X_u, X_v \in X$, $\Pr[I_S(X_u) \neq I_S(X_v)] \leq \alpha D \cdot \|X_u - X_v\|^2$, where I_S is the indicator function of S .

Bansal et al. [BFK⁺11] showed the existence of such separators (in the process also giving an efficient algorithm to construct them).

Lemma 3.3 ([BFK⁺11]). *For all $\beta < 1$ there exists an m -orthogonal separator with distortion $D = O\left(\sqrt{\frac{\log |X| \log m}{\beta}}\right)$.*

The dependency on β follows from calculations in Lemma 4.9 in [CMM06]. From the explanation of the above Lemma in [BFK⁺11], we know $\gamma = \sqrt{\beta}/8$, so the exponent in Lemma 4.9 in [CMM06] is $1/(1-\gamma^2) - 1 = O(\beta)$, and we want $(\log m'/m')^{O(\beta)}$ to be smaller than $1/m$. Setting $m' = m^{O(1/\beta)}$ suffices. Then the distortion is $O(\sqrt{\log |X| \log m'}) = O\left(\sqrt{\frac{\log |X| \log m}{\beta}}\right)$.

Now we show the following, which immediately implies Theorem 3.1.

Theorem 3.4. *For any $\delta > 0$, $0.25 > \beta > 0$, let $m = 10r^2/\delta$. Let D denote the best distortion possible for an m -orthogonal separator with separation β . If X is any set of vectors in ℓ_2^2 , one for each vertex in the graph, and the minimum expansion ϕ_{local} among subsets of at most $2n/r$ vertices satisfies $\phi_{\text{local}} \geq O(\phi_{SDP}D/\delta)$, then there exist r points S in X such that $\|X_S^\perp X\|_F^2 \leq O(\delta + \beta)\|X\|_F^2$.*

The actual construction of orthogonal separators from [CMM06] requires the origin to be inside the vector set. To achieve this, we will translate all vectors in the same direction:

Proposition 3.5. *If Theorem 2.1 fails, then there exists a set of vectors $X \in \ell_2^2$ with $0 \in X$.*

Proof. Given the vectors $[X_u]_u$ found by Theorem 2.1, we know that $\sum_u \|X_u\|^2 = \frac{1}{2} \mathbb{E}_u \sum_v \|X_u - X_v\|^2$. Hence there exists some t for which $\sum_u \|X_u - X_t\|^2 \leq 2 \sum_u \|X_u\|^2$. After having fixed such t , we define our new vectors as $X'_u \leftarrow X_u - X_t$. It is easy to see that $X' \in \ell_2^2$, $0 \in X'$ and for every subset of size $r - 1$, eq. (2) is satisfied (except ϵ becomes $\epsilon/2$, which only changes the constants in O notation). \square

We start by showing that most sets in the support of the orthogonal separator should be large.

Lemma 3.6. *If $\phi_{local} \geq 2\phi_{SDP}D/\delta$ as in the hypothesis of Theorem 3.4, and S is chosen according to the orthogonal separator, then $\mathbb{E}[|S| \cdot I_{|S| \leq 2n/r}] \leq \delta \mathbb{E}[|S|]$, where $I_{|S| \leq 2n/r}$ is the indicator variable for the event “ $|S| \leq 2n/r$.”*

Proof. On one hand we know

$$\mathbb{E}[\text{number of edges cut}] \geq \mathbb{E}[|S| \cdot I_{|S| \leq 2n/r}] \cdot \Phi_{local} \geq \mathbb{E}[|S| \cdot I_{|S| \leq 2n/r}] \cdot \phi_{local}/2.$$

On the other hand by condition 3 in the definition,

$$\mathbb{E}[\text{number of edges cut}] \leq \alpha D \sum C_{uv} \|X_u - X_v\|^2.$$

Since $\sum C_{uv} \|X_u - X_v\|^2 \leq \phi_{SDP} \sum \|X_u\|^2 = \phi_{SDP} \mathbb{E}[|S|]/\alpha$, we know

$$\mathbb{E}[|S| \cdot I_{|S| \leq 2n/r}] \leq \frac{1}{\phi_{local}} \alpha D \sum C_{uv} \|X_u - X_v\|^2 \leq \delta \mathbb{E}[|S|]. \quad \square$$

Now we state a corollary but first we need this definition.

Definition 3.7 (volume). The *volume* of a subset $X' \subset X$ is

$$vol(X') = \frac{\sum_{X_u \in X'} \|X_u\|^2}{\sum_{X_u \in X} \|X_u\|^2}.$$

Corollary 3.8. *There exists a subset $X' \subset X$ with volume at least $1 - 2\delta$, such that the following is true. Let S' be a set picked probabilistically by first picking S randomly according to the separator and letting S' be S if $|S| \geq 2n/r$, and the empty set otherwise. Then we have:*

1. For all $X_u \in X'$ we have $\Pr[X_u \in S'] \geq \alpha \|X_u\|^2/2$.
2. For all $X_u, X_v \in X$ with $\|X_u - X_v\|^2 \geq \beta \min\{\|X_u\|^2, \|X_v\|^2\}$,

$$\Pr[X_u \in S' \text{ and } X_v \in S'] \leq \frac{\min\{\Pr[X_u \in S], \Pr[X_v \in S]\}}{m}.$$

Proof. (Sketch) The first condition is by Markov. The second condition holds because the S' is always a subset of S , so the probability of LHS only decreases. \square

Now we are ready to prove Theorem 3.4.

Proof. (Theorem 3.4) We give an algorithm that shows iteratively picks r points such that most of the volume in X' lies close to them.

1. Initially none of the points are marked.
2. $i \leftarrow 1$.
3. While there is still a point in X' that is not marked:
 - (a) Let X_i be the point with largest norm among the unmarked points of X' .
 - (b) Pick a set S_i ($|S_i| \geq 2n/r$) containing X_i , and containing at most $2n/m$ points that have distance more than $\beta\|X_i\|^2$ from X_i . (Such a set exists as shown below.)
 - (c) Mark all points in S_i as well as all points that have distance at most $2\beta\|X_i\|^2$ from X_i . Denote by M_i the set of points that were previously unmarked and got marked in this step.
 - (d) Look over all M_j for $j < i$ and if any points in them have distance at most $\beta\|X_i\|^2$ to X_i then add them to M_i as well.
 - (e) $i \leftarrow i + 1$.

First we show using the probabilistic method why we can always perform step 3b. Pick a random set S' from the distribution of the separator, conditioning on its containing X_i . By the properties of S' we know if $\|X_i - X_v\|^2 \geq \beta\|X_i\|^2$, then the conditional probability $\Pr[X_v \in S' | X_i \in S'] \leq 2/m$. So the expected number of points in S' whose distance is at least $\beta\|X_i\|^2$ from X_i is at most $2n/m$, and in particular there must be one set that satisfy the condition.

Then we need to show that this process terminates in r steps. To do so it suffices to show that each S_i has at least n/r points that were not in any S_j for $j < i$. We know that $|S_i| > 2n/r$. We claim its intersection with any S_j for $j < i$ is at most $4n/m$. The reason is that X_i was unmarked at the start of this phase, which implies that the balls of radius $\beta\|X_j\|^2$ and $\beta\|X_i\|^2$ around X_j and X_i respectively must be disjoint (note that $\|X_j\| > \|X_i\|$) and thus the only intersections among S_i, S_j are from points outside these balls, which we know to be at most $4n/m$. Since $4nr/m < n/r$ (recall $m = 10r^2/\delta$), we have conclude that each S_i introduces at least n/r new points, so the process must terminate in r steps.

Finally we bound the average distance of the other points from this set $S = \{X_1, \dots, X_t\} (t \leq r)$, specifically, the quantity $\frac{\|X_S^\perp X\|_F^2}{\|X\|_F^2}$ by $O(\delta + \beta)$.

All points outside X' (the set in Corollary 3.8) anyway have volume at most 2δ , so their contribution is upperbounded by that. To bound the contribution of points in X' , consider how the sets M_1, \dots, M_t were picked. If X_u is $\beta\|X_i\|^2$ -close to X_i , then X_u is in M_i (these sets are disjoint by the construction). Otherwise X_u belongs to M_i where i is the time that X_u gets marked.

All points in M_i have norm at most $\|X_i\|^2$ since otherwise they would have been picked instead of X_i . Also more than n/r points (in fact, $2n/r - 2n/m > n/r$) in M_i are $\beta\|X_i\|^2$ -close to X_i , and at most $2n/m$ points are $2\beta\|X_i\|^2$ -far from X_i , so

$$\sum_{X_u \in M_i} \|X_u\|^2 \geq (|M_i| - 2n/m)(1 - 2\beta)\|X_i\|^2 \geq |M_i|\|X_i\|^2/3.$$

On the other hand, after projection to the space orthogonal to X_S , all but $2n/m$ points are smaller than $2\beta\|X_i\|^2$, therefore after projection

$$\sum_{X_u \in M_i} \|X_i^\perp X_u\|^2 \leq (|M_i| - 2n/m) \cdot 2\beta\|X_i\|^2 + 2n/m \cdot \|X_i\|^2 \leq O(\beta + \delta)|M_i|\|X_i\|^2.$$

Summing up the inequalities for all M_i 's we get the upperbound $O(\beta + \delta)$ needed for the theorem. \square

Algorithmic version. Since Bansal et al. [BFK⁺11] give an efficient algorithm for constructing orthogonal separators, the above proof immediately can be made algorithmic.

Corollary 3.9. *There is an algorithm that given a weighted graph $G = (V, E)$ in which $\phi_{\text{local}} > \frac{O(\sqrt{\log n \log r / \epsilon})}{\epsilon^{3/2}} \phi_{\text{sparsest}}$ computes a $(1 + \epsilon)$ -approximation to SPARSEST CUT in time $2^{O(r)} \text{poly}(n)$. Here ϕ_{local} is the minimum sparsity of sets of size at most $2n/r$.*

In fact the algorithm outputs one of the following.

1. *Either a subset with sparsity at most $(1 + \epsilon)\phi_{\text{SDP}}$,*
2. *Or a subset of size at most $\frac{2n}{r}$ with sparsity at most $\frac{O(\sqrt{\log n \log r / \epsilon})}{\epsilon^{3/2}} \phi_{\text{SDP}}$.*

Here ϕ_{SDP} is the optimum value of eq. (5) for $r + 3$ rounds.

Proof. (Sketch) Consider the algorithm from Theorem 2.1. If it outputs a partition, we are done. Otherwise, we apply the algorithm for constructing orthogonal separator in [BFK⁺11] on the set of vectors as constructed in Proposition 3.5. The above existence proof of the set S fails for this set of vectors, therefore Lemma 3.6 fails, and there must be a small set in the orthogonal separator that has desired expansion. \square

4 Bounded Genus Graphs

In this section, we prove an analog of our result for graphs with orientable genus g . The standard LP relaxation [LR99] for SPARSEST CUT on such graphs has an integrality gap of $O(\log g)$ [LS10]. For planar graphs (when $g = 0$), Park and Phillips [PP93] presented a weakly polynomial time algorithm for the problem of edge expansion using dynamic programming.

Here we show how to give a $(1 + \epsilon)$ -approximation when the graph satisfies a certain local expansion condition. Note that this expansion condition is true for instance in $O(1)$ -dimensional grids when $r = \text{poly}(1/\epsilon)$.

Theorem 4.1. *There is a polynomial-time algorithm that given a weighted graph G with orientable genus g in which $\phi_{\text{local}} > \frac{O(\log g)}{\epsilon^2} \phi_{\text{sparsest}}$ (where ϕ_{local} is the minimum sparsity of sets of size at most n/r) computes a $(1 + \epsilon)$ -approximation to SPARSEST CUT and similar problems in $2^{O(r)} \text{poly}(n)$.*

In fact the algorithm outputs one of the following.

1. *Either a subset with sparsity at most $(1 + \epsilon)\phi_{\text{SDP}}$,*
2. *Or a subset of size at most $\frac{n}{r}$ with sparsity at most $\frac{O(\log g)}{\epsilon^2} \phi_{\text{SDP}}$.*

Here ϕ_{SDP} is the optimum value of eq. (5) for $r + 2$ rounds.

Before proving Theorem 4.1, let us first recall the theory of random partitions of metric spaces, and its specialization to graphs of bounded genus. If (V, d) is a metric space then a *padded decomposition at scale Δ* is a distribution over partitions P of V where each block of P has diameter Δ . Its *padding parameter* is the smallest $\beta \geq 1$ such that the ball of radius Δ/β around a point has a good chance of lying entirely in the block containing the point:

$$\text{Prob}_P[B_d(u, \Delta/\beta) \subseteq P(u)] \geq \frac{1}{8} \text{ for all } u \in V. \quad (3)$$

The *padding parameter* of a graph G is the smallest β such that every semimetric formed by weighting the edges of G has a padded decomposition with padding parameter at most β . The following theorems are known.

Theorem 4.2. 1. [LS10] If G has orientable genus g , then its padding parameter is $O(\log g)$.

2. [FT03] If G has no $K_{p,p}$ minor, then its padding parameter is $O(p^2)$.

Our main technical lemma is the following.

Lemma 4.3. Given a graph $G = (V, E)$ and positive integer r , there exists an algorithm which runs in time $2^{O(r)} \text{poly}(n)$ and outputs one of the following for any $\epsilon > 0$:

1. Either a subset with sparsity at most $(1 + \epsilon)\phi_{\text{SDP}}$ where ϕ_{SDP} is the optimum value of eq. (5),

2. Or a subset of size at most $\frac{n}{r}$ with sparsity at most $\frac{O(\beta)}{\epsilon^2} \phi_{\text{SDP}}$.

Proof. The idea is to apply the algorithm from Theorem 2.1. If it finds a cut of sparsity $(1 + \epsilon)\phi_{\text{SDP}}$, then we are done. Otherwise let $[X_u]_u$ be the vectors output by it. We show how to use padded decompositions of the shortest-path semimetric given by distances $\|X_u - X_v\|^2$ and then produce a small nonexpanding set.

Let ν denote the average squared length of these vectors, i.e. $\nu \triangleq \frac{1}{n} \sum_u \|X_u\|^2$ so that $\nu = \mu(1 - \mu)$. Choose Δ at least $\frac{\epsilon}{2} \frac{\sum_u \|X_u\|^2}{n}$. Take a padded decomposition at scale Δ and pick a random partition P out of it.

CLAIM: The expected number of nodes that lie in subsets of size less than n/r in P is at least $\frac{\epsilon}{2} \sum_u \|X_u\|^2$.

Proof For each subset $S \in P$ with size $|S| \geq \frac{n}{r}$, if we choose an arbitrary $t \in S$, eq. (2) implies that:

$$\begin{aligned} \epsilon \sum_u \|X_u\|^2 &\leq \sum_{S \in P: |S| \geq \frac{n}{r}} \sum_{u \in S} \|X_t - X_u\|^2 + \sum_{T \in P: |T| < \frac{n}{r}} \sum_{u \in T} \|X_u\|^2 \\ &\leq \sum_{S \in P: |S| \geq \frac{n}{r}} \Delta |S| + \sum_{T \in P: |T| < \frac{n}{r}} \sum_{u \in T} \|X_u\|^2 \leq \frac{\epsilon}{2} \sum_u \|X_u\|^2 + \sum_{T \in P: |T| < \frac{n}{r}} |T|. \\ \frac{\epsilon}{2} \sum_u \|X_u\|^2 &\leq \sum_{T \in P: |T| < \frac{n}{r}} |T|. \end{aligned}$$

□

Now we choose a threshold $\tau \in [0, \Delta/\beta]$ uniformly at random. Then for each $T \in P$ with $|T| \leq \frac{n}{r}$, let $\hat{T} \subseteq T$ be the subset of nodes which are in the same partition block as the ball

of radius τ around them. We output such \hat{T} with minimum sparsity among all $T \in P$ with $|T| \leq \frac{n}{r}$. Using standard arguments, we can prove that any pair of nodes u and v is separated with probability at most $\frac{\|X_u - X_v\|^2}{\Delta/\beta}$. This means the total expected capacity cut will be at most $\frac{\beta}{\Delta} \sum_{u < v} C_{uv} \|X_u - X_v\|^2$. Moreover eq. (3) implies that:

$$\mathbb{E}_P \left[\sum_{T \in P: |T| \leq n/r} |\hat{T}| \right] \geq \frac{1}{8} \sum_{T \in P: |T| \leq n/r} |T| \geq \frac{\epsilon}{16} \sum_u \|X_u\|^2.$$

Putting all together, we see that there exists some $T \in P$ with $|T| \leq \frac{n}{r}$ such that

$$\phi_G(T) \leq \frac{O(\beta)}{\epsilon^2} \phi_{SDP}. \quad \square$$

Combining Lemma 4.3 with the bounds from Theorem 4.2 immediately implies Theorem 4.1.

5 Small-set Expander Flows

In [ARV09], *expander flows* are used as approximate certificate for expansion, which work for *all* values of expansion. (By contrast, the eigenvalue or spectral bound of Alon-Cheeger is most useful only for expansion close to $\Omega(1)$.) This section concerns *small-set expander flows* (SSE flows) which can be viewed as approximate certificates of the expansion of small sets. An (r, d, β) -SSE flow is a multicommodity flows in which small sets S (ie sets of size at most n/r for some small r) have $\beta d|S|$ outgoing flow where β is close to $\Omega(1)$. The flow is undirected, and the amount of flow originates at every node is at most d . Since the flow resides in the host graph and $\beta d|S|$ leaves every small set S , an (r, d, β) -SSE flow is trivially a certificate that small sets have edge expansion $\Omega(d\beta)$ in the host graph.

Of particular interest here will be a surprising connection between SSE flows and finding near-optimal SPARSEST CUT. In other words, information about expansion of small sets can be leveraged into knowledge about the expansion of all sets. We note that such a leveraging was already shown in [ABS10] using spectral techniques, but only when Small set expansion is $\Omega(1)$, roughly speaking (the reason is that the proof is Cheeger-like).

We note that given a flow it seems difficult (as far as we know) to verify that it is an SSE flow. Thus we will also be interested in a closely related notion of *spectral SSE flow*, which by contrast is easily recognized using eigenvalue computation. This is the one used in our algorithm.

Definition 5.1 (Spectral SSE Flow). A (r, d, λ) -*spectral SSE flow* is a multicommodity flow whose vertices have degree between $d/2$ and d , and the r^{th} smallest eigenvalue of its Laplacian matrix is at least $d\lambda$.

The relationship between the two types of flow rely upon the so-called higher order Cheeger inequalities [LRTV12, LGT12].

Theorem 5.2 (Rough statement). *If the graph has an (r, d, β) SSE flow then it also has an $(2r, d, \Omega(\beta^2/\log r))$ spectral SSE flow. Conversely, if the graph has an (r, d, λ) spectral SSE flow then it has a weaker version of $(r, d, \beta = \lambda)$ combinatorial SSE flow.*

See Lemma B.22 and Lemma B.24 for more precise statements.

Now we describe how these results are useful. First, just *existence* of SSE flows is enough to imply a low integrality gap for the Lasserre relaxation. This is reminiscent of primal-dual frameworks (e.g., expander flows being a family of dual solutions for the ARV SDP relaxation and thus giving a lower bound on the optimum) but we don't know how to make that formal yet.

Theorem 5.3. *If a (r, d, λ) -spectral SSE flow exists in the graph for $d\lambda \gg \frac{1}{\epsilon}\phi_{\text{sparsest}}$, then the GS rounding algorithm computes a $(1 + \epsilon)$ -approximation to SPARSEST CUT when applied on the $O(r/\epsilon)$ -level Lasserre solution. In particular, the integrality gap of the Lasserre relaxation is at most $(1 + \epsilon)$.*

The other result is a more direct approximation algorithm that does not use SDP hierarchies at all. Instead it uses a form of spectral rounding (as in [ABS10]) that produces a set with low symmetric difference to the optimum sparsest cut, followed by the clever idea of Andersen and Lang [AL08] to purify this set into a bonafide cut of low expansion.

Theorem 5.4. *There is a $2^{O(r)}\text{poly}(n)$ time algorithm that given a graph and a (r, d, λ) -spectral SSE flow for $d\lambda \gg \frac{1}{2}\phi_{\text{sparsest}}$ outputs a cut of sparsity at most $(1 + \epsilon)\phi_{\text{sparsest}}$.*

The above two theorems become important only because of the following two theorems which concern the *existence* of the flow.

Theorem 5.5. *If $d \ll \Phi_{\text{local}}\sqrt{\log r}/\sqrt{\log n}$ then the graph has a $(r, d, \Omega((\log r)^{-2}))$ SSE flow.*

This theorem follows from Lemma B.19 and Lemma B.20.

Theorem 5.6. *If $d \ll \Phi_{\text{local}}\sqrt{\log r}/\sqrt{\log n}$ then the graph has a $(2r, d, \Omega((\log r)^{-5}))$ spectral SSE flow. Furthermore, a $(4r, d, \Omega((\log r)^{-5}))$ spectral SSE flow can be found in polynomial time.*

This theorem follows Theorems 5.2 and 5.5 and Lemma B.25. The algorithm to find the spectral SSE flow uses the fact that maximizing the sum of first r eigenvalues of a matrix is a convex objective.

In fact, when Φ_{local} is small, we can actually find a small set that does not expand well.

Theorem 5.7. *For any graph $G = (V, E)$ and any value d , there is a polynomial time algorithm that either finds a $(4r, d, \Omega((\log r)^{-5}))$ spectral SSE flow, or finds a set of size at most $100n/r$ that has expansion at most $O(d\sqrt{\log n}/\sqrt{\log r})$.*

For more details see Lemma B.26.

5.1 Overview of proof of existence of SSE flows

To keep the main paper relatively concise, we have move the proof of existence to the appendix and give an overview here.

From a distance, the existence proof for SSE flows uses similar ideas as the one for expander flows in [ARV09]: we write an exponential size LP that is feasible iff the desired flow exists, and then reason about the properties of dual solutions (using properties of flows, cuts, and ℓ_2^2 metrics) to show that the LP is feasible.

We write an LP that enforces each vertex has degree at most d in the flow, and for every set S of size $n/3r$ to n/r , the amount of outgoing flow is at least $\beta d|S|$, the precise LP can be found in Appendix B.1.

The dual of this LP consists of a nonnegative weight s_i for all vertices and w_e for each edge, and also a nonnegative weight for every set of size between $n/3r$ to n/r . We shall prove the following Lemma:

Lemma 5.8 (imprecise). *Given a valid dual solution with degree d and β parameter $= \Theta((\log r)^{-2})$, there is an algorithm that finds a set of size at most $100n/r$ with expansion $O(d\sqrt{\log n}/\sqrt{\log r})$.*

In order for the algorithm to run in polynomial time, we first need to represent the LP dual concisely, and as stated above it involves a nonnegative weight on exponentially many cuts! As in ARV, this concise representation is possible since a nonnegative weighting of cuts is an ℓ_1 metric and the algorithm is only interested in the “distance” between two vertices in this metric (which is the measure of sets that contains one of the vertices but not the other). The ℓ_1 metric can be concisely represented by some ℓ_2^2 vectors; see Appendix B.2.1 and Lemma B.26.

The proof of the Lemma above uses the “chaining” idea from [ARV09], but there are many differences which we list here.

- (a) The proof is handicapped since it is only allowed to use *local expansion* (i.e., expansion of sets of size at most $O(n/r)$), and this requires us to invent novel ways of applying the region-growing framework in [LR99] (see Appendix B.2.3). Many steps in our algorithms rely on such region growing arguments, including Lemmas B.6, B.16 and B.18.
- (b) In [ARV09] all vectors have unit norm, here however the ℓ_2^2 vectors can have different norms. We use a known reduction that transforms the vectors for a large subset of vertices, so that they are in a sphere of fixed radius. See Appendix B.2.4.
- (c) The existence of matching covers used in the ARV proof is unclear and has to be carefully established. This uses a certain “spreading constraint” that holds for ℓ_2^2 metrics supported on small sets. See Lemma B.15. Also, a matching cover may not exist because a set of vertices is far away from other vertices in graph distance (distance according to the weights on edges). We call such sets *obstacle sets of type I*, and use region-growing arguments to remove these sets, see Appendix B.2.5.
- (d) The crux of the ARV proof is to prove the existence of a special pair of vertices that are close in graph metric (i.e., the metric given by the weights on the edges) and far apart in ℓ_2^2 metric). From the existence proof and global expansion Φ_{global} , one can immediately establish the existence of $\Omega(n)$ such pairs, which is needed in the argument. The analogous idea does not work here since the proof is handicapped by being restricted to only use local expansion. However, we show that this step can only fail if there exists an *obstacle set of type II*. We design another region-growing type argument to handle this; see Lemmas B.12, B.16 and B.18.
- (e) The ARV argument uses Alon-Cheeger inequality: for d regular graphs, the second eigenvalue of the Laplacian is $\Omega(1)$ iff the graph has expansion $\Omega(1)$. The analogous result for small set expansion, the so-called “higher order Cheeger inequality,” has only recently been established, and only in one direction and in a weaker form [LRTV12, LGT12]. This weak form makes us lose extra $\text{poly}(\log r)$ factors in many theorems which are potentially improveable. For details see Lemma B.22.

5.2 Finding Sparsest Cut using SSE flow

Before we delve into the long proof of existence of SSE flows, we quickly show how they are useful in approximating SPARSEST CUT. As mentioned, there are two methods for this.

5.2.1 Method 1: Using Lasserre Hierarchy Relaxation

This will use a modification of an idea of Guruswami-Sinop which we now recall. Recall (see Appendix A) that the solutions for $r' + 2$ rounds of Lasserre Hierarchy relaxation satisfies the following property:

$$\frac{\sum_{u < v} C_{uv} \|X_u - X_v\|^2}{\sum_u \|X_u\|^2} = \frac{\text{Tr}(X^T X L(G))}{\|X\|_F^2} = \phi_{SDP},$$

where the approximation ratio is bounded by $(1 - \frac{\|X_S^\perp X\|_F^2}{\|X\|_F^2})^{-1}$ over all sets S of size r' by Theorem 2.1.

Theorem 5.9 (Theorem 3.2 in [GS13]). *Given positive integer $r \geq 0$ and positive real $\epsilon > 0$, the above approximation ratio is upperbounded by $\left(1 - \frac{1}{1-\epsilon} \frac{\sum_{i>r} \sigma_i(X^T X)}{\|X\|_F^2}\right)^{-1}$ for $r' = \frac{r}{\epsilon} + r + 1$.*

Proof. (Sketch) Using the column based low-rank matrix reconstruction error bound from [GS12b], it can be shown that there exists set S of size $r' = r/\epsilon + r - 1$ such that the numerator $\|X_S^\perp X\|_F^2 \leq (1 - \epsilon)^{-1} \sum_{j \geq r+1} \sigma_j(X^T X)$, where $\sigma_j(X^T X)$ is the j^{th} largest eigenvalue of $X^T X$. \square

In order to bound the sum of eigenvalues, the analysis in [GS13] uses von Neumann's trace inequality, which we present in a slightly more general form:

Proposition 5.10. *For any matrix $Y \succeq 0$ and positive integer r :*

$$\sum_{i \geq r+1} \sigma_i(Y) = \min_{Z \succeq 0} \frac{\text{Tr}(Y \cdot Z)}{\lambda_{r+1}(Z)}.$$

In the original analysis of [GS13], this claim is used with $Y \leftarrow X^T X$ and $Z \leftarrow L(G)$, whereupon one obtains:

$$\frac{\sum_{i>r} \sigma_i(X^T X)}{\|X\|_F^2} \leq \frac{\text{Tr}(X^T X \cdot L(G))}{\lambda_{r+1}(G) \|X\|_F^2} \leq \frac{\phi_{SDP}}{\lambda_{r+1}(G)}.$$

Thus

$$\frac{\sum_{i \leq r} \sigma_i(X^T X)}{\|X\|_F^2} \geq 1 - \frac{\phi_{SDP}}{\lambda_{r+1}(G)}. \quad (4)$$

Consequently, the rounding analysis in [GS13] requires a bound on the λ_{r+1} value of the graph.

Our idea is to use Proposition 5.10 by substituting the Laplacian of the spectral SSE flow as Z in the above calculation, and then use the lowerbound on the λ_r value of this flow Laplacian. This uses the following easy lemma.

Lemma 5.11. *If X is described above, then for any flow F that lies in the host graph G :*

$$\frac{\sum_{i>r} \sigma_i(X^T X)}{\|X\|_F^2} \leq \frac{\phi_{SDP}}{\lambda_{r+1}(F)}.$$

Proof. Since F is routable in G and $X \in \ell_2^2$, we have:

$$\text{Tr}(X^T X \cdot L(F)) \leq \text{Tr}(X^T X \cdot L(G)) \leq \phi_{SDP} \|X\|_F^2.$$

Choosing $Y \leftarrow X^T X$ and $Z \leftarrow L(F)$, we see that the Claim implies:

$$\frac{\sum_{i>r} \sigma_i(X^T X)}{\|X\|_F^2} \leq \frac{\text{Tr}(X^T X \cdot L(F))}{\lambda_{r+1}(F) \|X\|_F^2} \leq \frac{\phi_{SDP}}{\lambda_{r+1}(F)}. \quad \square$$

Now Theorem 5.3 follows using Lemma 5.11 and eq. (4).

Remark: Note that we only need $\lambda_{r+1}(F)$ to be more than ϕ_{SDP} . Such flows could potentially exist under more general conditions than our local expansion condition.

5.2.2 Method 2: Using Subspace Enumeration and Cut Improvement

We show that given a (r, d, λ) spectral SSE flow, where $d\lambda$ is much larger than the expansion Φ of sparsest cut, it is possible to use eigenspace enumeration idea of [ABS10] together with the ideas of [AL08] to get a good approximation to SPARSEST CUT.

Lemma 5.12 (Eigenspace Enumeration, [ABS10]). *There is a $2^{O(r)} n^{O(1)}$ time algorithm that, given a graph whose r^{th} smallest eigenvalue (of Laplacian) is $\lambda_r \geq 20\Phi/\epsilon$, outputs a set of subsets $X \subset \{0, 1\}^V$ with the following guarantee: for every subset S that has expansion Φ , there is a vector $x \in X$ such that*

$$\frac{|x - \vec{1}_S|}{|\vec{1}_S|} \leq \frac{8\Phi}{\lambda_r}.$$

The above eigenspace enumeration allows us to compute a “guess” that has low symmetric difference with the optimum cut. Then we can use a simple version of cut improvement algorithm of [AL08] to improve it to a cut of low expansion.

Lemma 5.13. *There is a $2^{O(r)} n^{O(1)}$ time algorithm that given a graph $G = (V, E)$, and a (r, d, λ) spectral SSE flow embeddable in G , enumerates $2^{O(r)} n^{O(1)}$ sets with the following guarantee. For any set S of size at most $n/2$ that has expansion $\Phi(S) \ll d\lambda\epsilon\delta$ (for $\epsilon + \delta < 1$), there is a set T in the output such that $\frac{|T \Delta S|}{|S|} \leq \delta$ and $\Phi(T) \leq (1 + \epsilon)\Phi(S)$ (Δ denotes symmetric difference).*

Proof. The capacity of flow that crosses S in the spectral SSE flow can only be smaller than $\Phi(S) \cdot |S|$ because the flow is embeddable in G . Hence when we apply Lemma 5.12 on the flow, we know there is a vector $\vec{1}_T$ in X that is $\epsilon\delta/2$ close to the indicator vector of S .

Using this vector, suppose we know the expansion $\Phi(S)$ (later we shall see we only need to know this value up to multiplicative factor, so the algorithm will enumerate all possible values). Construct a single commodity flow instance where we add a source s and sink t to the graph. For each vertex $i \in T$, there is an edge from i to sink t with capacity $4\Phi(S)/\delta$. For each vertex $i \notin T$, there is an edge from source s to i with capacity $4\Phi(S)/\delta$.

Now we find the min-cut that separates source s and sink t . Since T is close to S , we know the capacity of this cut is at most $(1 + \epsilon/2)\Phi(S)|S|$ because $\{s\} \cup S$ achieves this capacity. Let the vertices that are on the same side with sink be Q , then we know $|Q \Delta T| \leq \frac{(1 + \epsilon/2)\Phi(S)|S|}{4\Phi(S)/\delta} \leq |S|\delta/2$.

Therefore $\frac{|Q \Delta S|}{|S|} \leq \frac{|Q \Delta T| + |T \Delta S|}{|S|} \leq \delta$.

On the other hand, the expansion of Q is at most

$$\frac{(1 + \epsilon/2)\Phi(S)|S| - |Q\Delta T| \cdot 4\Phi(S)/\delta}{|S| - |S\Delta T| - |Q\Delta T|} = \frac{(1 + \epsilon/2)\Phi(S) - 4x\Phi(S)/\delta}{(1 - \epsilon\delta/2) - x} \leq (1 + \epsilon)\Phi(S).$$

(where in the second step, we substituted $x \triangleq \frac{|Q\Delta T|}{|S|}$). \square

Corollary 5.14. *Given graph $G = (V, E)$ and a (r, d, λ) spectral SSE flow embeddable in G . There is a $2^{O(r)}n^{O(1)}$ time algorithm that:*

- Finds a set S with $\phi(S) \leq (1 + O(\epsilon))\phi_{\text{sparsest}}$ if $d\lambda \gg \phi_{\text{sparsest}}/\epsilon^2$;
- Finds a set S with $\Phi(S) \leq (1 + O(\epsilon))\Phi_{\text{global}}$ if $d\lambda \gg \Phi_{\text{global}}/\epsilon$;
- Finds a set S of size at least $cn/2$ such that $\Phi(S) \leq (1 + O(\epsilon))\Phi_{c\text{-balanced}}$ if $d\lambda \gg \Phi_{c\text{-balanced}}/c\epsilon$.

Proof. (sketch) For sparsest cut, choose $\delta = \epsilon$ in Lemma 5.13. For edge expansion, choose $\delta = 1/2$. For c -balanced separator, choose $\delta = c/2$. \square

6 Conclusions

The fact that it is possible to compute $(1 + \epsilon)$ -approximation for SPARSEST CUT on an interesting family of graphs seems very surprising to us. Further study of Guruswami-Sinop rounding also seems promising: our analysis is still not using the full power of their theorem.

Our work naturally leads us to the following imprecise conjecture, which if true would yield immediate progress.

Conjecture: (Imprecise) *In “interesting” families of graphs —ie those where existing algorithms for SPARSEST CUT fail— $\Phi_{\text{local}}/\Phi_{\text{global}}$ is large, say $\gg \sqrt{\log n}$.*

As support for this conjecture we observe that if our algorithm does not beat $\sqrt{\log n}$ -approximation on some graph, then there is a constant r and a set of size n/r whose expansion is at least $\sqrt{\log n}$ times the optimum.

Furthermore, it is conceivable that SSE flows exist in graphs even when the local expansion condition is not met. For our analysis of the rounding algorithm from [GS13] we only need the existence of an SSE flow of degree say $> 1.1\phi_{\text{sparsest}}$ (see Section 5.2.1). Conceivably such flows exist in a wider family of graphs, and this could be another avenue for progress.

Acknowledgements

We gratefully acknowledge helpful discussions with Venkat Guruswami, Jon Kelner, Ravi Krishnaswamy, Assaf Naor and David Steurer. This work was funded by grants from the NSF and the Simons Foundation.

References

- [ABS10] Sanjeev Arora, Boaz Barak, and David Steurer. Subexponential algorithms for Unique Games and related problems. In *FOCS*, pages 563–572, 2010.

- [AHK04] Sanjeev Arora, Elad Hazan, and Satyen Kale. $o(\sqrt{\log n})$ approximation to sparsest cut in $\tilde{O}(n^2)$ time. In *FOCS*, pages 238–247, 2004.
- [AK07] Sanjeev Arora and Satyen Kale. A combinatorial, primal-dual approach to semidefinite programs. In *STOC*, pages 227–236, 2007.
- [AL08] Reid Andersen and Kevin J. Lang. An algorithm for improving graph partitions. In *SODA*, pages 651–660, 2008.
- [ALN08] Sanjeev Arora, James Lee, and Assaf Naor. Euclidean distortion and the sparsest cut. *J. American Mathematical Society*, 21(1):1–21, 2008.
- [AM85] Noga Alon and V. D. Milman. λ_1 , isoperimetric inequalities for graphs, and superconcentrators. *J. Comb. Theory, Ser. B*, 38(1):73–88, 1985.
- [ARV09] Sanjeev Arora, Satish Rao, and Umesh V. Vazirani. Expander flows, geometric embeddings and graph partitioning. *J. ACM*, 56(2), 2009.
- [BFK⁺11] Nikhil Bansal, Uriel Feige, Robert Krauthgamer, Konstantin Makarychev, Viswanath Nagarajan, Joseph Naor, and Roy Schwartz. Min-max graph partitioning and small set expansion. In *FOCS*, pages 17–26, 2011.
- [BHK⁺12] Boaz Barak, Aram Harrow, Jonathan Kelner, David Steurer, and Yuan Zhou. Hypercontractivity, Sum-of-Squares proofs, and their applications. In *STOC*, pages 307–326, 2012.
- [BRS11] Boaz Barak, Prasad Raghavendra, and David Steurer. Rounding semidefinite programming hierarchies via global correlation. In *FOCS*, pages 472–481, 2011.
- [CMM06] Eden Chlamtac, Konstantin Makarychev, and Yury Makarychev. How to play unique games using embeddings. In *FOCS*, pages 687–696, 2006.
- [DKSV06] Nikhil R. Devanur, Subhash Khot, Rishi Saket, and Nisheeth K. Vishnoi. Integrality gaps for sparsest cut and minimum linear arrangement problems. In *STOC*, pages 537–546, 2006.
- [DR10] Amit Deshpande and Luis Rademacher. Efficient volume sampling for row/column subset selection. In *FOCS*, pages 329–338, 2010.
- [FT03] Jittat Fakcharoenphol and Kunal Talwar. An improved decomposition theorem for graphs excluding a fixed minor. In *RANDOM-APPROX*, pages 36–46, 2003.
- [GS11] Venkatesan Guruswami and Ali Kemal Sinop. Lasserre hierarchy, higher eigenvalues, and approximation schemes for graph partitioning and quadratic integer programming with PSD objectives. In *FOCS*, pages 482–491, 2011.
- [GS12a] Venkatesan Guruswami and Ali Kemal Sinop. Faster SDP hierarchy solvers for local rounding algorithms. In *FOCS*, pages 197–206, 2012.
- [GS12b] Venkatesan Guruswami and Ali Kemal Sinop. Optimal column-based low-rank matrix reconstruction. In *SODA*, pages 1207–1214, 2012.

- [GS13] Venkatesan Guruswami and Ali Kemal Sinop. Approximating non-uniform sparsest cut via generalized spectra. In *SODA*, 2013.
- [GT13] Shayan Oveis Gharan and Luca Trevisan. Improved ARV Rounding in Small-set Expanders and Graphs of Bounded Threshold Rank. *ArXiv e-prints*, April 2013.
- [KLL⁺13] Tsz Chiu Kwok, Lap Chi Lau, Yin Tat Lee, Shayan Oveis Gharan, and Luca Trevisan. Analysis of spectral partitioning through higher order spectral gap. In *STOC*, 2013.
- [Las02] Jean B. Lasserre. An explicit equivalent positive semidefinite program for nonlinear 0-1 programs. *SIAM J. Optimization*, 12(3):756–769, 2002.
- [LGT12] James R. Lee, Shayan Oveis Gharan, and Luca Trevisan. Multi-way spectral partitioning and higher-order cheeger inequalities. In *STOC*, pages 1117–1130, 2012.
- [LR99] Frank Thomson Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, 1999.
- [LRTV12] Anand Louis, Prasad Raghavendra, Prasad Tetali, and Santosh Vempala. Many sparse cuts via higher eigenvalues. In *STOC*, pages 1131–1140, 2012.
- [LS10] James R. Lee and Anastasios Sidiropoulos. Genus and the geometry of the cut graph. In *SODA*, pages 193–201, 2010.
- [LS11] James R. Lee and Anastasios Sidiropoulos. Near-optimal distortion bounds for embedding doubling spaces into ℓ_1 . In *STOC*, pages 765–772, 2011.
- [MMV12] Konstantin Makarychev, Yury Makarychev, and Aravindan Vijayaraghavan. Approximation algorithms for semi-random partitioning problems. In *STOC*, pages 367–384, 2012.
- [MN04] Manor Mendel and Assaf Naor. Euclidean quotients of finite metric spaces. *Advances in Mathematics*, 189(2):451–494, 2004.
- [OZ13] Ryan O’Donnell and Yuan Zhou. Approximability and proof complexity. In *SODA*, 2013.
- [Par03] Pablo A. Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Math. Program.*, 96(2):293–320, 2003.
- [PP93] James K. Park and Cynthia A. Phillips. Finding minimum-quotient cuts in planar graphs. In *STOC*, pages 766–775, 1993.
- [RS10] Prasad Raghavendra and David Steurer. Graph expansion and the unique games conjecture. In *STOC*, pages 755–764, 2010.
- [She09] Jonah Sherman. Breaking the multicommodity flow barrier for $o(\sqrt{\log n})$ -approximations to sparsest cut. In *FOCS*, pages 363–372, 2009.

A Overview of Lasserre Hierarchy Relaxation for Sparsest Cut and Rounding

In this section, we will give a brief description of Lasserre Hierarchy relaxation for Uniform Sparsest Cut problem, the rounding algorithm of [GS13] and its analysis.

We present the formal definitions of Lasserre Hierarchy relaxations [Las02], tailored to the setting of the problems we are interested in, where the goal is to assign to each node in V a label from $\{0, 1\}$.

Definition A.1 (Lasserre vector set). Given a set of variables V and a positive integer r , a collection of vectors x is said to satisfy r -rounds of Lasserre Hierarchy, denoted by $x \in \text{Lasserre}_r(V)$, if it satisfies the following conditions:

1. For each set $S \in \binom{V}{\leq r+1}$, there exists a function $x_S : \{0, 1\}^S \rightarrow \mathbb{R}^\Upsilon$ that associates a vector of some finite dimension Υ with each possible labeling of S . We use $x_S(f)$ to denote the vector associated with the labeling $f \in \{0, 1\}^S$. For singletons $u \in V$, we will use x_u and $x_u(1)$ interchangeably. For $f \in \{0, 1\}^S$ and $v \in S$, we use $f(v)$ as the label v receives from f . Also given sets S with labeling $f \in \{0, 1\}^S$ and T with labeling $g \in \{0, 1\}^T$ such that f and g agree on $S \cap T$, we use $f \circ g$ to denote the labeling of $S \cup T$ consistent with f and g : If $u \in S$, $(f \circ g)(u) = f(u)$ and vice versa.
2. $\|x_\emptyset\|^2 = 1$.
3. $\langle x_S(f), x_T(g) \rangle = 0$ if there exists $u \in S \cap T$ such that $f(u) \neq g(u)$.
4. $\langle x_S(f), x_T(g) \rangle = \langle x_A(f'), x_B(g') \rangle$ if $S \cup T = A \cup B$ and $f \circ g = f' \circ g'$.
5. For any $u \in V$, $\sum_{j \in \{0, 1\}} \|x_u(j)\|^2 = \|x_\emptyset\|^2$.
6. (implied by above constraints) For any $S \in \binom{V}{\leq r+1}$, $u \in S$ and $f \in \{0, 1\}^{S \setminus \{u\}}$, $\sum_{g \in \{0, 1\}^u} x_S(f \circ g) = x_{S \setminus \{u\}}(f)$.

One can view $\|x_S(f)\|^2$ as the “probability” of f , in which case the corresponding “conditional” probabilities are given by $\frac{\langle x_S(f), x_u \rangle}{\|x_S(f)\|^2}$. Our relaxation is the following:

$$\begin{aligned} \min_{\mu, x} \quad & \frac{1}{n\mu(1-\mu)} \sum_{u < v} C_{u,v} \|x_u - x_v\|^2 \\ \text{st} \quad & \frac{1}{n} \sum_u x_u = \mu x_\emptyset, \\ & \|x_\emptyset\|^2 = 1, \quad x \in \text{Lasserre}_{r'+2}(V), \quad \mu \in \{1/n, 2/n, \dots, 1/2\}. \end{aligned} \tag{5}$$

Note that we can easily eliminate the variable μ from eq. (5) by enumerating over all $\frac{n}{2}$ values. For the special case of uniform sparsest cut, the rounding algorithm from [GS13] can be summarized as follows. Given a feasible solution of eq. (5):

1. Let $X_u \triangleq x_u - \frac{1}{n} \sum_v x_v = x_u - \mu x_\emptyset$.
Observe (i) $\|X_u - X_v\|^2 = \|x_u - x_v\|^2$; (ii) $\sum_u X_u = 0$; (iii) $\sum_u \|X_u\|^2 = \frac{1}{n} \sum_{u < v} \|x_u - x_v\|^2 = \mu(1 - \mu)$; (iv) $1 - \mu \geq \|X_u\|^2 \geq \mu^2$.
2. Choose a set S of r' nodes using column selection [GS12b, DR10] from $[X_u]_u$.

3. Sample $f : S \rightarrow \{0, 1\}$ with probability proportional to $\|x_S(f)\|^2$.
4. Perform threshold rounding using the “conditional probabilities” assigned for each node $u \in V$ which is proportional to $\langle x_S(f), x_u \rangle$.

B Constructing SSE Flows

B.1 Definition and LP Formulation of SSE Flows

Recall for any graph $G = (V, E)$, a multicommodity flow in G assigns demand $\delta_{i,j}$ for pairs of vertices i, j , and simultaneously route $\delta_{i,j}$ units of flow from i to j for all pairs while satisfying capacity constraints. In particular, if we use f_p to denote the amount of flow routed along path p , $\mathcal{P}_{i,j}$ to denote all paths from i to j , then a multicommodity flow should satisfy the following constraints:

$$\forall i, j \in V \quad \sum_{p \in \mathcal{P}_{i,j}} f_p = \delta_{i,j} \quad (6)$$

$$\forall e \in E \quad \sum_{p \in \mathcal{P}} f_p \leq c_e \quad (7)$$

We shall only consider symmetric flows (i.e. $\delta_{i,j} = \delta_{j,i}$). For a multicommodity flow, we call $\delta_i = \sum_{j \in V} \delta_{i,j}$ the degree of vertex i . Now we can define SSE flows:

Definition B.1 (SSE flow). A (r, d, β) -SSE flow is a multicommodity flow whose vertices have degree at most d , and for any set S of size at most n/r , we have

$$\sum_{i \in S, j \notin S} \delta_{i,j} \geq \beta d |S|.$$

For a flow, we will define the expansion of a set S to be $\frac{\sum_{i \in S, j \notin S} \delta_{i,j}}{d|S|}$, so the requirement of SSE flow is just the expansion of all small sets should be at least β .

SSE flow is a complicated object, we will also use two weaker versions of SSE flow. The first one is especially useful for the LP formulation:

Definition B.2 (Weak SSE flow). A (r, d, β) weak SSE flow is a multicommodity flow whose vertices have degree at most d , and for any set S of size between $n/3r$ and n/r , we have

$$\sum_{i \in S, j \notin S} \delta_{i,j} \geq \beta d |S|.$$

Notice that the idea of restricting set S to have roughly size n/r is also used in [ARV09] (where in the LP formulation the sets have size $n/6$ to $n/2$). We will use the LP formulation for weak SSE flows:

$$\forall i \in V \quad \sum_j \sum_{p \in \mathcal{P}_{i,j}} f_p \leq d \quad (8)$$

$$\forall e \in E \quad \sum_{p \in \mathcal{P}} f_p \leq c_e \quad (9)$$

$$\forall S \subset V, n/3r \leq |S| \leq n/r \quad \sum_{i \in S, j \notin S} \sum_{p \in \mathcal{P}_{i,j}} f_p \geq \beta d |S| \quad (10)$$

Later we show this LP is feasible for some values of β and d by showing that the dual LP is not feasible. The dual LP is given by

$$\sum_e c_e w_e + d \sum_{i \in V} s_i < \beta d \sum_S z_S |S| \quad (11)$$

$$\forall i, j, p \in \mathcal{P}_{i,j} \quad \sum_{e \in p} w_e + s_i + s_j \geq \sum_{S: i \in S, j \notin S} z_S \quad (12)$$

$$z_S, w_e, s_i \geq 0 \quad (13)$$

Here s_i ($i \in V$), w_e ($e \in V$) and z_S ($S \subset V, n/3r \leq |S| \leq n/r$) are the dual variables corresponding to eqs. (8), (9) and (10) respectively. Since the dual LP is homogeneous, all variables can be scaled simultaneously without affecting the validity of a solution. Throughout this section we shall always assume the following normalization:

$$\sum_S z_S |S| = n \quad (14)$$

We would like SSE flows to serve as approximate certificate of small set expansion. However, small set expansion is in general hard to certify, even if the expansion is close to 1. Hence we give a weaker (but still useful) form, spectral SSE flow, which can be easily certified, and is closely related to combinatorial SSE flows by high order Cheeger's inequality[LRTV12, LGT12].

Definition B.3. A (r, d, λ) spectral SSE flow is a multicommodity flow whose vertices have degree between $d/2$ and d , and the r^{th} smallest eigenvalues of the Laplacian of the graph is at least $d\lambda$.

B.2 Existence of SSE Flows

The main result of this Section is the existence of weak SSE flows.

Lemma B.4. *For any graph $G = (V, E)$ and any $d > 0$ either there is a set of size at most $100n/r$ that has expansion smaller $d\sqrt{\log n}/\sqrt{\log r}$, or there exists a $(r, d, \Omega(\log^{-2}r))$ weak SSE flow embeddable in G .*

B.2.1 ℓ_2^2 Mapping of z_S

In order to show that weak SSE flow exist, we argue that the dual LP does not have any valid solution. In fact, we show something even stronger: given any dual solution that satisfies eq. (11), eq. (13) and eq. (14), there is a polynomial time algorithm that either finds a nonexpanding set of size at most $100n/r$, or finds a path where eq. (12) is violated.

The dual solution has exponentially many variables. We shall use a compressed description that is good enough for the algorithm: all the variables z_S are mapped to n vectors Z_1, \dots, Z_n , such that for all i, j , $\|Z_i - Z_j\|_2^2 = \sum_{i \in S, j \notin S} z_S$. This is possible because Z_i 's can have one coordinate for each set S , if i is in S then the coordinate is $\sqrt{z_S}$, otherwise the coordinate is just 0.

The upper-bound on the size of S implies for any i , there are at most Cn/r points within ℓ_2^2 distance $\|Z_i\|_2^2/C$ for all $C > 1$. This is because for any set S of Cn/r points, if we pick a random point j in S , with probability 1 the expected distance between i and j is at least $\mathbb{E}_{j \in S}[\|Z_i - Z_j\|_2^2] \geq \sum_{t \ni i} z_S \Pr[j \notin t] \geq \sum_{t \ni i} z_S / C = \|Z_i\|_2^2 / C$. To avoid giving out exponentially many variables, the dual solution will only contain Z_i 's that satisfy this property.

We shall rewrite the constraints for Z_i 's

$$\sum_e c_e w_e + d \sum_{i \in V} s_i < \beta d n \quad (15)$$

$$z_S, w_e, s_i \geq 0 \quad (16)$$

$$\sum_{i \in V} \|Z_i\|_2^2 = n \quad (17)$$

$$\forall i \in V, C > 1 \quad |\{j : \|Z_j - Z_i\|_2^2 \leq 4 \|Z_i\|_2^2 / 5\}| \leq 5n/4r \quad (18)$$

$$\forall i \in V \quad \|Z_i\|_2^2 \leq 3r \quad (19)$$

$$\forall i, j \in V, p \in \mathcal{P}_{i,j} \quad \sum_{e \in p} w_e + s_i + s_j \geq \|Z_i - Z_j\|_2^2 \quad (20)$$

The last constraint is called the *spreading constraint*. A candidate dual solution is just a set of variables $s_i (i \in V)$, $w_e (e \in E)$, $Z_i (i \in S)$ that satisfies these constraints.

Consider w_e 's as edge distances on the graph, and let $d_{i,j}$ be the shortest path distance between i and j with weights w_e . From now on we refer to this weighted distance as the *graph distance*. Intuitively, given a candidate dual solution, the algorithm either finds a nonexpanding set or finds two vertices who have small s_i 's, small distance in graph distance and large distance in ℓ_2^2 distance.

B.2.2 Proof Idea

Given the dual solution, we try to apply arguments similar to [ARV09] in order to find a pair of vertices that are close in graph distance, but far in ℓ_2^2 metric. When this pair of vertices also have small s_i 's, it violates Equation (20) hence contradicting the feasibility of dual solution.

In [ARV09] this proof goes by projecting all points along a random direction and arguing that there must be many pairs of points that are close in graph distance but far in projection distance: they call it a *matching cover*. However, constructing a matching cover in our setting case is highly nontrivial, because the proof is only allowed to use local expansion. We adapt the region-growing argument in [LR99] in novel ways to solve this problem, see Appendix B.2.3.

The first difficulty in the argument is that each vertex might have very different $\|Z_i\|_2$ (that is, they are in very different measure of sets in the dual solution). But this is easily fixed by embedding the points into a single scale using ideas from [ALN08] (see Lemma B.6 in Appendix B.2.4).

It turns out that in order for a matching cover to not exist, one of the following two types of *obstacles* must exist, detailed discussion appears in Appendix B.2.5.

The first type of obstacle set is a set whose D_0 -neighborhood in graph distance (D_0 is a parameter that will be chosen later) contains only $O(n/r)$ points. Intuitively, this is an obstacle since it would be hard to match these vertices to other vertices within graph distance D_0 because they simply don't have enough neighbours. We show that the total volume of such sets cannot be too large using the region-growing framework, see Lemma B.11.

The second type of obstacle set is a set of at most $10n/r$ vertices with large s_i whose D_0 neighbourhood in graph distance contains only $O(n/r)$ points with small s_i . Intuitively such sets are bad because we want to construct matching covers only on vertices with small s_i (in order to get the final contradiction with Equation (20)). Such sets would mean it is possible for a set S with small s_i to be only close to vertices with large s_i 's, and it would be impossible to match all

the vertices in S with vertices with small s_i 's. We again use region-growing arguments to remove such sets. The number of vertices removed cannot be large, because otherwise the sum of s_i 's will be too large and violates eq. (11) (see Lemma B.12).

Without these obstacle sets, it becomes possible to construct a matching cover (see Lemma B.16 in Appendix B.2.6). This matching cover allows us to adapt arguments in [ARV09] (see Lemmas B.17 and B.18), and conclude that either there is a short path (in graph distance) that crosses many cuts, or there is a nonexpanding set. The first case contradicts with the validity of the dual solution. In the second case we get a nonexpanding set, which again implies the existence of obstacle sets of type I or II.

B.2.3 Region-growing Argument

As mentioned earlier, a key component of our proof is the region-growing argument from [LR99]. This argument applies to an undirected graph whose edges have arbitrary nonnegative capacities. The goal (in [LR99]) is to give a partition into blocks that have low *diameter* (distance being measured using edge weights) and on average have *few* edges crossing between the blocks. The tradeoff between these two quantities is controlled by the expansion of the underlying unweighted graph. Here we view this argument as giving an efficient *partition oracle*, which maintains a set of vertices V_i at step i ($V_0 = V$). At step i the oracle takes a set $S_i \subset V_{i-1}$ of size at least $n/F(r)$ where F is a fixed polynomial, and then outputs S'_i , $S_i \subset S'_i \subset V_{i-1}$, and updates $V_i = V_{i-1} \setminus S'_i$. There is a “center” $j \in S_i$ such that every other $j' \in S_i$ has distance at most D_0 to j (we specify D_0 later in Lemma B.5).

At step t , we say the partition maintained by the oracle is the collection of disjoint sets $S'_1, S'_2, \dots, S'_t, V_t$. The capacity of edges in the partition is always at most $n\alpha/20\Delta \log 30r$.

Lemma B.5 ([LR99]). *Let $G = (V, E)$ be a graph with edge capacities c_e and edge lengths w_e . Let W be the total weighted edge length: $W = \sum_{e \in E} c_e w_e$. Then for any polynomial $F(r)$, and any $D_0 = C\Delta \log 30r \cdot \log r W / n\alpha$ (where C is a constant depending on F), there is an efficient partition oracle whose partitions always have capacity at most $n\alpha/20\Delta \log 30r$.*

Proof. The proof is similar to Lemma 3 in [LR99]. However there the region-growing procedure starts from a single vertex (and the loss is $\log n$ because n is roughly the ratio between the volume of the graph and the volume of a single vertex). Here instead we start region-growing from the sets given to the oracle. Because the sets all have large volume (more than $n/\text{poly}(r)$), we lose only a $\log r$ factor. \square

B.2.4 Reducing to Single Scale

The region growing argument applies to a particular scale Δ . However, not all vertices have ℓ_2^2 norm close to that scale. In this part we show how to reduce the problem to a single scale Δ .

Lemma B.6. *Given a dual solution with $\beta \leq C_\beta(\log r)^{-3/2}$, there is an algorithm that finds a Δ , and calls a partition oracle with scale Δ and size $F(r)$. After the algorithm, the number of remaining vertices in the oracle is at least $\frac{n}{5\Delta \log 30r}$, and all but $n/F(r)$ of the remaining vertices satisfy one of the two properties:*

1. $\|Z_i\|_2^2 \geq \Delta/2$.

2. $s_i \geq D_2/10 = \Omega(\Delta/\log r)$.

The value D_2 comes from Lemma B.18 and will be $\Omega(\Delta/\log r)$.

First we use an averaging argument to find Δ .

Lemma B.7. *There exists some threshold $0.1 < \Delta < 3r$ such that the number of vertices with $\|Z_i\|_2^2 \geq \Delta$ is at least $\frac{n}{4\Delta \log 30r}$.*

Proof. We shall bucket the vertices according to $\|Z_i\|_2^2$. There will be $b = \lceil \log 30r \rceil$ buckets, the u -th ($u \in \{1, \dots, b\}$) bucket B_u contains vertices with $\|Z_i\|_2^2$ in range $[0.1 * 2^{u-1}, 0.1 * 2^u)$. There will be one extra bucket B_0 which contains vertices with $\|Z_i\|_2^2$ in range $[0, 0.1)$. By Equation (19), $\|Z_i\|_2^2 \leq 3r$ so these buckets cover all vertices.

We know $\sum \|Z_i\|_2^2 = n$, let $l_u = \sum_{i \in B_u} \|Z_i\|_2^2$, then $\sum_{u=0}^b l_u = n$. We also know $l_0 \leq 0.1n$, so there must be a bucket B_u with $l_u \geq 0.9n/b$. Choose $\Delta = 0.1 * 2^{u-1}$, we know the number of vertices with $\|Z_i\|_2^2 \geq \Delta$ is at least the size of B_u , which is at least $l_u/2\Delta \geq n/4\Delta b$. \square

Now we are ready to prove Lemma B.6.

Proof of Lemma B.6. Take the value Δ from Lemma B.7. Let Q be the set of vertices whose $\|Z_i\|_2^2$ is at most $\Delta/2$ and s_i is at most $D_2/10$. Let B be the set of vertices whose $\|Z_i\|_2^2$ is at least Δ .

If the size of Q is at most $n/F(r)$, then the Lemma is true. Otherwise, use the partition oracle to separate the set Q . From the oracle we get a Q' which contains everything in Q . Consider any vertex $i \in Q' \cap B$, by definition of oracle we know there is a vertex $j \in Q$ such that $D_{i,j} \leq D_0$. On the other hand, $\|Z_i - Z_j\|_2^2 \geq \|Z_i\|_2^2 - \|Z_j\|_2^2 \geq \Delta/2$. By eq. (20) we know $s_i > \Omega(\Delta)$, since $\sum_{i \in V} s_i \leq \beta n$, the size of $Q' \cap B$ is at most $\beta n/\Delta < |B|/50$. The size of current set of the oracle is at least $n/5\Delta \log 30r$. \square

We shall also use the following Lemma from [ALN08, MN04] to project everything to a ball of squared radius Δ .

Lemma B.8 ([MN04]). *There exists a mapping $T : \ell_2 \rightarrow \ell_2$ such that $\|T(z)\|_2 \leq \sqrt{\Delta}$ for all $z \in \ell_2$ and for all $z, z' \in \ell_2$*

$$\frac{1}{2} \leq \frac{\|T(z) - T(z')\|_2}{\min\{\sqrt{\Delta}, \|z - z'\|_2\}} \leq 1.$$

As a corollary, we now prove the following.

Corollary B.9. *There is a mapping that maps Z_i to X_i , such that $\|X_i\|_2 = \sqrt{\Delta/2}$ for all $i \in V$, and for all $i, j \in V$*

$$\frac{1}{8} \leq \frac{\|X_i - X_j\|_2^2}{\min\{\Delta, \|Z_i - Z_j\|_2^2\}} \leq 1.$$

Proof. Just let $X_i = \frac{1}{\sqrt{2}}T(Z_i) \oplus (\sqrt{\Delta - \|T(Z_i)\|_2^2})$ where \oplus denotes concatenation of vectors. It is easy to verify the claim. \square

After mapping all the Z_i 's to X_i 's, for a vertex i with $\|Z_i\|_2^2 \geq \Delta/2$, vertices that are within squared distance $\Delta/20$ in X metric are also within squared distance $2\Delta/5$ in Z metric (Corollary B.9). By spreading constraints there can only be at most $5n/4r$ such vertices.

B.2.5 Obstacle Sets

The plan of the proof is to apply cover composition from [ARV09] in order to find a short path (in graph metric) that crosses a lot of cuts. At any step i , let V_i be the remaining vertices in the partition oracle. Let Q_i be the set of vertices in V_i that have large s values (at least $D_2/10$ as in Lemma B.6. In this case there are two kinds of obstacle sets that prevents us from applying the cover composition argument.

Definition B.10 (Obstacle Sets). At some step i of the partition oracle, a set $S \subset V_i$ is an obstacle set of type I if it has size at least $n/F(r)$, and the D_0 neighbourhood contains at most $100n/r$ vertices in V_i .

A set $T \subset Q_i$ is an obstacle set of type II, if it has size at least $10n/r$, and the D_0 neighbourhood contains at most $100n/r$ vertices in $V_i \setminus Q_i$.

Using region-growing arguments, we can remove the obstacle sets using the partition oracle without removing many vertices.

Lemma B.11. *For a partition oracle with distance D_0 as in Lemma B.6, if at some step i , $\sum_{j \leq i, |S'_j| \leq 100n/r} |S'_j| \geq n\alpha/10\Delta \log 30r$, then one of the S'_j of size at most $100n/r$ has expansion at most α .*

In particular, there is a set $H \subset V_i$ whose size is at least $|V_i| - n/10\Delta \log 30r$, such that any subset $S \subset H$ of size at least $n/F(r)$ expands to at least $100n/r$ vertices in H .

Proof. If we take the sum of capacity of all outgoing edges from these S'_j , each edge in the partition is counted at most twice, therefore

$$\sum_{j \leq i, |S'_j| \leq 100n/r} |E(S'_j, V \setminus S'_j)| \leq n\alpha/10\Delta \log 30r.$$

On the other hand we know the sum of sizes is at least $n/10\Delta \log 30r$, by averaging argument there must be one set that has expansion α . \square

For proving Lemma B.19 this α will be chosen as $O(d\sqrt{\log n}/\sqrt{\log r})$.

Notice that it is very important that the algorithm always uses the partition oracle when it wants to remove a set of vertices. If the algorithm simply removes a set of vertices, it will be hard to bound the number of edges cut, and Lemma B.11 is no longer true. In this case we may have many obstacle sets of type I and cannot find a matching cover.

For obstacle sets of type II, since a large fraction of the vertices in their neighbourhood have large s_i , they cannot cover a lot of vertices without contradicting the validity of the dual solution

Lemma B.12. *Use the partition oracle in Lemma B.6 to remove obstacle sets of type II. At any step, let II be the set of steps where a set of type II is removed. Then $\sum_{j \in II} |S'_j| \leq n/20\Delta \log 30r$.*

Proof. By the definition of obstacle sets of type II, we know each S'_j contains at least $10n/r$ vertices with s -value at least $D_2/10$. On the other hand, it contains at most $90n/r$ vertices with s -value smaller than $D_2/10$. Therefore $1/10$ fraction of the vertices in S'_j have s value at least $D_2/10$.

$$\sum_{u \in V} s_u \geq \sum_{j \in II} \sum_{u \in S'_j} s_u \geq \frac{|S'_j|}{10} \cdot \frac{D_2}{10}.$$

On the other hand $\sum_{u \in V} s_u \leq \beta n$, so when $\beta = C \log^{-2} r$ for small enough C we know $\sum_{j \in II} |S'_j| \leq n/20\Delta \log 30r$. \square

B.2.6 Gaussian Projections and Matching Covers

Recall the definitions of Matching Covers and Uniform Matching Covers in [ARV09]:

Definition B.13. A (σ, δ, c') -matching cover of a set of points is a set \mathcal{M} of matchings such that for at least a fraction δ of directions u , there exists a matching $M_u \in \mathcal{M}$ of at least $c'n$ pairs of points, such that each pair $(i, j) \in M_u$ are within graph distance $2D_0$, and satisfies

$$\langle X_i - X_j, u \rangle \geq 2\sigma\sqrt{\Delta}/\sqrt{d}.$$

The associated matching graph M is defined as the multigraph consisting of the unions of all matchings M_u .

Definition B.14. A set of matchings \mathcal{M} (σ, δ) -uniform-matching-covers a set of points S if for every unit vector u , there is a matching M_u of S such that every $(i, j) \in M_u$ is $2D_0$ close in graph distance, satisfies $|\langle u, X_i - X_j \rangle| \geq 2\sigma\sqrt{\Delta}/\sqrt{d}$, and for every i , $\mu(u : i \text{ matched in } M_u) \geq \delta$.

Notice that in addition to the properties in [ARV09], we further require that every matched pair must be close in graph distance.

Let the dimension of X_i 's be d . Let u be a uniformly random unit vector, when d is large enough we know

Lemma B.15. *There exists thresholds $0 < \theta_1 < \theta_2$ such that $\theta_2 - \theta_1 = \Omega(\sqrt{\log r}/\sqrt{d})$ and polynomial $G(r)$. Let $\text{Good}(u)$ be the event that number of vertices with projection more than $\theta_1\Delta$ is smaller than $5n/r$. For any vertex i whose $\|Z_i\|_2^2 \geq \Delta/2$, $\Pr[u \cdot X_i \geq \theta_2\sqrt{\Delta} \text{ and } \text{Good}(u)] \geq 1/G(r)$.*

Proof. (sketch) We know each such i has at most $5n/4r$ closeby points. For points that are not close, conditioned on i has large projection, the probability that they also have pretty large projection is very small. Hence conditioned on i being in, the expected number of vertices that have large projection is small. By Markov we know $\Pr[\text{Good}(u) | u \cdot X_i \geq \theta_2\sqrt{\Delta}] \geq 1/2$. \square

Given a set of vertices V_t , which can be partitioned into three parts P, Q, R , vertices $i \in P$ all have $\|Z_i\|_2^2 \geq \Delta/2$, vertices $j \in Q$ all have $s_j \leq D_2/10$, and $|R| \leq n/F(r)$ (notice that this is exactly what's guaranteed by Lemma B.6), we will use the following algorithm to find matching covers:

CONSTRUCT COVER (P,Q,R)

1. Pick uniformly random unit vector u .
2. Let $Left = \{i : i \in P \text{ and } \langle X_i, u \rangle \geq \theta_2\sqrt{\Delta}\}$, $Right = \{i : i \in P \text{ and } \langle X_i, u \rangle \leq \theta_1\sqrt{\Delta}\}$.
3. While exists pair $i \in Left$ and $j \in Right$ within graph distance $2D_0$.
4. Match (i, j) , remove i, j from $Left, Right$.
5. Fail if $|P \setminus Right| < 5n/r$ and number of unmatched vertices in $Left$ is at least $n/F(r)$

If the algorithm fails, the following Lemma shows that we will have an obstacle set of type I or II.

Lemma B.16. *If Construct Cover fails, then it finds an obstacle set of I or II.*

Proof. Let S be the set of vertices that are left unmatched in *Left*. Let $\Gamma_{D_0}(S)$ be the D_0 neighbourhood of S in $P \cup Q \cup R$. If $|\Gamma_{D_0}(S)| \leq 100n/r$ first case of the Lemma is satisfied.

If $|\Gamma_{D_0}(S)| > 100n/r$, then either $|\Gamma_{D_0}(S) \cap P| > 80n/r$, in which case by simple counting argument there must be a point left in *Right* that is close to some point in S , and these two vertices can be matched (this contradicts with the assumption). When $|\Gamma_{D_0}(S) \cap P| \leq 90n/r$, let $T = \Gamma_{D_0}(S) \cap Q$. Clearly $|T| > 10n/r$, and $\Gamma_{D_0}(T) \cap P \subset \Gamma_{2D_0}(S) \cap P$. The number of $2D_0$ neighbours of S in P cannot be more than $90n/r$ (otherwise we will be able to find a matching pair), hence $|\Gamma_{D_0}(T) \cap P| \leq 90n/r$ and $|\Gamma_{D_0}(T) \cap (P \cup R)| \leq 90n/r + n/F(r) < 100n/r$. \square

[ARV09] has a Lemma that shows matching covers imply uniform matching covers. However in our situation, in order to apply the cover composition Lemma, we need a really large uniform matching cover, which is not guaranteed by the Lemma in [ARV09].

If Construct Cover does not fail with polynomial probability, then Lemma B.15 means the matching cover is already “almost” uniform, in the sense that if we ignore the fact that $n/F(r)$ points will not be matched, each vertex will be in the matching with probability at least $1/G(r)$.

Lemma B.17. *If Construct Cover fails with probability less than $1/n^2G(r)$, then there is a set $W \subset P$ of size at least $|P| - 4|P|/r$ that is $(1/G(r)r, \theta_2 - \theta_1)$ uniformly matching covered.*

Proof. Consider the matching graph. First, even for the $n/F(r)$ points that remains unmatched, consider that they are matched to something. In this case each vertex has degree at least $1/G(r)$ by Lemma B.15.

Now remove the edges that correspond to unmatched edges. In this step we have removed at most $|P| \cdot 1/G(r)r$ volume. Then we repeatedly remove any vertex that has degree at most $1/G(r)r$. Again we will remove at most $|P|/G(r)r$ volume. So the total volume removed is bounded by $2|P|/G(r)r$.

However, we know that each vertex in P started with degree at least $1/G(r)$. Removing $2|P|/G(r)r$ volume can reduce the degree of at most $4|P|/r$ vertices to below $1/2G(r)$. Therefore at most $4|P|/r$ vertices are removed. \square

B.2.7 Adapting ARV

Using the uniform matching cover constructed above, and mechanisms in [ARV09], we can get the following Lemma.

Lemma B.18. *If $W \subset V$ and has $(1/G(r)r, \Omega(\sqrt{\log r}/\sqrt{d}))$ uniform matching cover. Then there exists an algorithm that either finds $i, j \in W$, such that $d_{i,j} \leq D_1 = O(D_0 \cdot \sqrt{\log n}/\sqrt{\log r})$, and $\|X_i - X_j\|^2 \geq D_2 = \Omega(\Delta/\log r)$, or finds a set whose $2D_0$ neighbourhood has size smaller than $100n/r$ in P .*

Proof. The proof follows from [ARV09], the algorithm basically follows the cover composition proof, maintaining the cover S_k along the induction steps (this is possible because the probabilities we are dealing with are all larger than some inverse polynomial, and the probabilities do not need to be estimated exactly). The main differences are:

1. Here we need to boost the probability from $1/G(r)r$ to $1 - 1/r$, this is $D_2 = \Omega(\Delta/\log r)$ (in [ARV09] we can find a pair that are constant distance away in ℓ_2^2 metric).

2. The definition of non-expanding set is now a set that does not expand to $100n/r$ vertices within graph distance $2D_0$. This is OK because either there is a pair within graph distance $2D_0$ and ℓ_2^2 distance more than D_2 , in which case the Lemma is true; or all vertices in this neighbouring set are also close in ℓ_2^2 distance, which then matches the definition of non-expanding set in [ARV09]. \square

B.2.8 Final Proof

The following Lemma immediately implies Lemma B.4.

Lemma B.19. *Given a dual solution with degree d and expansion $\beta < C_\beta \log^{-2} r$ (where C_β is a universal constant), there is an algorithm that finds a set of size at most $100n/r$ with expansion $O(d\sqrt{\log n}/\sqrt{\log r})$.*

Proof. First apply Lemma B.6. If Lemma B.6 did not find a set, then we have sets P, Q, R from Lemma B.6 and a partition oracle whose current set is $P \cup Q \cup R$.

Now we shall repeatedly apply Construct Cover. In this case we can get an obstacle set of type I or II.

If it fails with more than $1/n^2 G(r)$ probability then we get an obstacle set from Lemma B.16. Otherwise we would have a large uniform matching cover by Lemma B.17. Then we apply Lemma B.18 on this uniform matching cover, under the assumptions $D_1 < D_2/10$ ⁵, since $W \subset P$ the first case of Lemma B.18 cannot happen. We must get a non-expanding set. Lemma B.16 also applies to this non-expanding set and we can again get an obstacle set of type I or II.

Once we get the obstacle set, feed that set into the partition oracle, and recurse on the current set of the oracle. We always call obstacle sets of type I S , and obstacle sets of type II T . The corresponding sets returned by the oracle will be called S' and T' , respectively.

At the end one of the two cases will happen: Either the sets corresponding to S' take up more than $n/10\Delta \log 30r$ vertices or the sets corresponding to T' take up more than $n/10\Delta \log 30r$ vertices.

In the first case Lemma B.11 shows one of the S' must have low expansion.

The second case contradicts the feasibility of dual solution because of Lemma B.12. \square

B.3 Getting SSE Flows and Spectral SSE Flows

Thus far our existence proof dealt with weak SSE flows.

B.3.1 Getting to SSE flows

Lemma B.20. *If $G = (V, E)$ is a graph and a (r, d, β) weak SSE flow is embeddable in G , either there is a set of size at most n/r that has expansion less than $d\beta$, or there exists a $(r, d, \beta/6)$ SSE flow embeddable in G .*

Proof. Let F be the weak SSE flow. If for all sets S of size $|S| \leq n/3r$, the F has expansion at least β , then F is already a SSE flow.

When there exists S of size smaller than $n/3r$ and the expansion in F is smaller than β , remove S (for remaining vertices replace edges going to S with self-loops) and repeat this procedure.

⁵Notice that the constant in D_1 is in fact hiding in the expansion $O(d\sqrt{\log n}/\sqrt{\log r})$ which can be chosen independently of D_2 .

If the union of the removed sets is U , the size of U cannot be larger than $n/3r$: if after removing some S the size of U first become larger than $n/3r$, then since S has size smaller than $n/3r$, the size of U must be between $n/3r$ and $2n/3r$. The expansion of U is at most the maximum expansion among sets S , which is smaller than β . Such a set cannot exist by the definition of weak SSE flows.

Now add a source and a sink to the graph. Add an edge from source to every vertex in U with capacity $d\beta$, add an edge from every vertex in $V \setminus U$ to the sink with capacity $d\beta$, and then try to route the maximum single-commodity flow from source to sink.

If the maximum flow is smaller than $d\beta|U|$, then there must be a cut of value smaller than $d\beta|U|$ in the new graph. Let Q be one side of this cut that contains the source, then $E(Q, V \setminus Q) < d\beta|U| - |Q \oplus U|d\beta$ (this is because, for every i in Q but not U , it has degree $d\beta$ to the sink; for every i in U but not Q , it has degree $d\beta$ from the source), and $|Q| \geq |U| - |Q \oplus U|$. The expansion of Q is strictly smaller than $d\beta$.

If the maximum flow is $d\beta|U|$, let the single-commodity flow be F_1 , and let $F_2 = (F + F_1)/2$ (here “+” just take the linear combination of demands). Clearly F_2 is still embeddable into G . For any set S of size at most n/r , if more than $|S|/3$ of the vertices are outside U , then it already has $\beta d|S|/6$ outgoing edges outside U in $F/2$; if less than $|S|/3$ of the vertices are outside U , then it has $d\beta|S|/6$ outgoing edges just by the flow $F_1/2$. Therefore F_2 is a $(r, d, \beta/6)$ SSE flow. \square

Unfortunately, this Lemma is only existential. In general, even if we are given a SSE flow, it is hard to verify it exactly.

B.3.2 Getting Spectral SSE flow

We can use higher order equivalents of Cheeger’s Inequality to establish a relation between SSE flows and spectral SSE flows:

Theorem B.21 ([LRTV12, LGT12]). *For any graph G , $\Phi_r \leq O(\sqrt{\lambda_{2r}(\mathcal{L}) \log r})$. Here $\lambda_{2r}(\mathcal{L})$ is the $2r$ -th smallest eigenvalue of the normalized Laplacian of G .*

This implies that if the largest and smallest degree are close, then an SSE flow is already a spectral SSE flow.

Lemma B.22. *For any graph $G = (V, E)$, if there is a (r, d, β) SSE-flow embeddable in G , then there is a $(2r, d, \Omega(\beta^2 / \log r))$ spectral SSE flow embeddable in G .*

Before proving Lemma B.22, we will need the following simple claim so as to relate the eigenvalues of *normalized* Laplacian matrix to the original Laplacian.

Claim B.23. *Let d_{\min}, d_{\max} be the minimum and maximum degrees in G , respectively. Then:*

$$\frac{1}{d_{\max}} L(G) \preceq \mathcal{L}(G) \preceq \frac{1}{d_{\min}} L(G).$$

Proof. For any pair of nodes u, v , $d_{\min} \leq \sqrt{d_u d_v} \leq d_{\max}$. Hence for any $x \in \mathbb{R}^V$:

$$\frac{x^T L(G) x}{d_{\max}} = \sum_{u < v} \frac{C_{uv}}{d_{\max}} (x_u - x_v)^2 \leq x^T \mathcal{L} x = \sum_{u < v} \frac{C_{uv}}{\sqrt{d_u d_v}} (x_u - x_v)^2 \leq \frac{x^T L(G) x}{d_{\min}}. \quad \square$$

Proof of Lemma B.22. Let F be the (r, d, β) SSE-flow, let F_1 be a flow whose demands are $\delta_{i,j} = c_{i,j}$. Clearly F_1 is embeddable in G and has degree 1. Let $F_2 = F/2 + dF_1/2$, then the degrees of vertices in F_2 are between $d/2$ and d .

By definition of SSE flow we know $\Phi_r(F_2) \geq \beta/2$. Let \mathcal{L} be the normalized Laplacian of F_2 , and L be its Laplacian, then by Theorem B.21 $\lambda_{2r}(\mathcal{L}) \geq \Omega(\beta^2/\log r)$.

Since the degrees of F_2 are all between $d/2$ and d , by Claim B.23, the eigenvalues of its normalized Laplacian are closely related to its Laplacian: $\lambda_{2r}(L) \geq \frac{d}{2}\lambda_{2r}(\mathcal{L}) = \Omega(d\beta^2/\log r)$. \square

The inverse direction (spectral flows imply combinatorial flows) is also true, except the combinatorial expansion must be defined on r disjoint sets instead of one set.

Lemma B.24 ([KLL⁺13]). *A (r, d, λ) spectral flow satisfies the following combinatorial expansion property: for any r disjoint sets S_1, S_2, \dots, S_r , the maximum of the expansion of these sets is at least $\lambda/2$.*

Proof. This proof comes from [KLL⁺13], we restate it here for completeness. We use Courant-Fischer-Weyl characterization the variational definition of r^{th} smallest eigenvalue:

$$\lambda_r(L) = \min_{\text{subspace } P \text{ of dimension } r} \max_{h \in P} \frac{h^T L h}{h^T h}.$$

Let the subspace P be the span of the indicator vectors of S_i 's. For any $h = \sum_{i=1}^r \lambda_i \vec{1}_{S_i}$, for all $u, v \in V$,

$$(h(u) - h(v))^2 \leq \sum_{i=1}^r 2\lambda_i^2 (\vec{1}_{S_i}(u) - \vec{1}_{S_i}(v))^2$$

So the Rayleigh Quotient of h is at most

$$R(h) = \max_{h \in P} h^T L h / \|h\|_2^2 \leq \frac{2 \sum_{i=1}^r \lambda_i (\vec{1}_{S_i}(u) - \vec{1}_{S_i}(v))^2}{\sum_{i=1}^r \lambda_i^2 \|\vec{1}_{S_i}\|_2^2} \leq 2 \max_{i \in [r]} R(\vec{1}_{S_i}).$$

We know $\max R(h) \geq \lambda$, so the maximum expansion must be at least $\lambda/2$. \square

In order to find spectral SSE flows, the following algorithm uses a convex program:

Lemma B.25. *If there exists a (r, d, λ) spectral SSE flow embeddable in G , there is an efficient algorithm that finds a $(2r, d, \lambda/2)$ spectral SSE flow.*

Proof. The algorithm tries to solve the following optimization problem:

$$\begin{aligned} & \max \sum_{i=1}^{2r} \lambda_i(L(F)) \\ & s.t. \forall i \in V \quad \frac{d}{2} \leq \sum_{j \in V} \sum_{p \in \mathcal{P}_{i,j}} f_p \leq d. \\ & F \text{ embeddable in } G. \end{aligned}$$

Here $L(F)$ is the Laplacian of the flow. The first constraint just says the degree of every vertex should be between $d/2$ and d . This is a convex program because entries of $L(F)$ are linear functions over f_p , and the sum of first $2r$ eigenvalues of a matrix is a concave function. The convex program can be solved in polynomial time.⁶

⁶There are exponentially many paths, but there is a canonical way of reducing the number of variables for flows.

Clearly the (r, d, λ) spectral SSE flow is a feasible solution and has objective value at least $rd\lambda$. Hence the solution of this convex program must have objective function at least $rd\lambda$, which means the $2r$ -th eigenvalue of $L(F)$ is at least $\frac{rd\lambda}{2r} = d\lambda/2$. \square

B.4 Finding a Small Nonexpanding Set when Eigenspace Enumeration Fails

Combining Theorems 5.4 and 5.6, we know if Φ_{local} for a graph is at least $O(\Phi_{global}\sqrt{\log n}\log^{4.5} r/\epsilon)$, there is an eigenspace enumeration algorithm that finds a $(1 + \epsilon)$ approximation of sparsest cut. Here we show when the algorithm fails, how to find a small set that does not expand in polynomial time.

Lemma B.26. *Given a graph G , for any d, r , there is a polynomial time algorithm that either finds a $(2r, d, \lambda = \Omega((\log r)^{-5}))$ spectral flow, or finds a set of size at most $100n/r$ that has expansion at most $O(d\sqrt{\log n}/\sqrt{\log r})$.*

Proof. By Lemmas B.20 and B.22, we know a weak SSE flow implies a spectral SSE flow unless there is a small set with very small expansion. Therefore if the algorithm in Lemma B.25 does not work, either there is a set of size at most $n/3r$ that has expansion $d\beta$ where $\beta = \theta((\log r)^{-2})$, or there is no weak SSE flow.

In the first case we can simply run the approximation algorithm for small set expansion in [BFK⁺11], which gives a $\sqrt{\log n \log r}$ approximation, the set we get will be small and has expansion at most $d\beta\sqrt{\log n \log r} < O(d\sqrt{\log n}/\log r)$.

In the second case, there is no weak SSE flow, so the LP for the weak SSE flow must be infeasible, and its dual must be feasible. The original dual formulation has exponentially many variables, however in Appendix B.2.1 we mapped the solution to a concise representation using l_2^2 vectors Z_i 's. Equations (15) to (20) are almost constraints of a semidefinite program, except for eq. (18). However, we can write the spreading constraint in more tractable way:

$$\forall i \in V \quad \sum_{j \in V} \min\{0.9 \|Z_i\|_2^2, \|Z_j - Z_i\|_2^2\} \geq 0.9 \|Z_i\|_2^2 n - \|Z_i\|_2^2 \cdot \frac{n}{r}.$$

This equation is clearly satisfied by the Z_i 's converted from the original dual solution, because there the number of vectors within $0.9 \|Z_i\|_2^2$ is at most $10n/9r$, even if all of them are identical with Z_i , the sum on the LHS can only be $\|Z_i\|_2^2 \cdot \frac{n}{r}$ away from its maximum possible value $0.9 \|Z_i\|_2^2 n$.

On the other hand, if this equation is satisfied, we know for any i , the number of j such that $\|Z_j - Z_i\|_2^2 \leq 0.8 \|Z_i\|_2^2$ is at most $10n/r$. This is very similar to Constraint (18) except the constants are larger. This increase in constant does not change anything in the proof of Lemma B.19.

Therefore, we can solve the SDP to get a concise representation of the dual solution, and then apply Lemma B.19 to find a set that has size at most $100n/r$ with expansion $O(d\sqrt{\log n}/\sqrt{\log r})$. \square

C Planted Expander Model

Our algorithm naturally applies to the *planted expander model*. In this model the graph has a planted bisection of expansion $\Phi_{planted}$, the smaller side of the bisection has size ρn . The induced subgraph on each side of the partition is an expander with expansion $\Phi_{global} \gg \Phi_{planted}\sqrt{\log n \log 1/\rho}$. In this case we can show the assumptions in Theorem 3.1 hold, and the algorithm gives a good approximation to sparsest cut.

This result is similar to the “planted spectral expander model” in [MMV12]. The main difference is that they assume the induced graphs of the partition have *algebraic* expansion constant times more than Φ_{planted} . Notice that our combinatorial expansion property only implies algebraic expansion of $\Phi_{\text{planted}}^2 \log n \log 1/\rho$, which might be smaller than Φ_{planted} if the planted bisection is sparse enough. Unfortunately our result only applies to regular graphs, therefore a comparison is not possible *per se*. Our formal guarantee is given in the following theorem.

Theorem C.1. *Assume graph $G = (V, E)$ is a regular graph with an unknown planted bisection $(S, V \setminus S)$. The size of S is ρn ($\rho \leq 1/2$) with $\Phi(S) = \Phi_{\text{planted}}$. If the induced subgraphs of S and $V \setminus S$ both have expansion $\Phi \gg \frac{1}{\epsilon^{1.5}} \Phi(S) \sqrt{\log n \log \frac{1}{\rho \epsilon}}$, then the algorithm in Theorem 3.1 with $r = O(1/\rho)$ gives a $(1 + \epsilon)$ approximation to sparsest cut.*

Proof. We only need to show the assumptions in Theorem 3.1 are satisfied: Sets of size $\rho n/2$ should have sparsity at least $\Omega(\phi_{\text{sparsest}} \sqrt{\log n \log \frac{1}{\rho \epsilon}} \epsilon^{-1.5})$. Since sparsity and expansion are within a constant factor, we will show the expansion of small sets are at least $\Delta \triangleq \Omega(\Phi(S) \sqrt{\log n \log \frac{1}{\rho \epsilon}} / \epsilon^{1.5})$.

For any set T of size at most $\rho n/2$, let T_1 be $T \cap S$ and T_2 be $T \cap (V \setminus S)$. By assumption we know $E(T_1, S \setminus T_1) \gg |T_1| \cdot \Delta$ and $E(T_2, V \setminus (S \cup T_2)) \gg |T_2| \Delta$. Hence

$$\Phi(T) = \frac{E(T, V \setminus T)}{|T|} \geq \frac{E(T_1, S \setminus T_1) + E(T_2, V \setminus (S \cup T_2))}{|T_1| + |T_2|} \gg \Delta. \quad \square$$